arXiv:1402.2127v2 [cs.LO] 11 Feb 2014

Deciding Entailments in Inductive Separation Logic with Tree Automata

Radu Iosif, Adam Rogalewicz, and Tomáš Vojnar

CNRS/Verimag, France and FIT BUT, Czech Republic
radu.iosif@imag.fr, {rogalew,vojnar}@fit.vutbr.cz

Abstract. Separation Logic (SL) with inductive definitions is a natural formalism for specifying complex recursive data structures, used in compositional verification of programs manipulating such structures. The key ingredient of any automated verification procedure based on SL is the decidability of the entailment problem. In this work, we reduce the entailment problem for a non-trivial subset of SL describing trees (and beyond) to the language inclusion of tree automata (TA). Our reduction provides tight complexity bounds for the problem and shows that entailment in our fragment is EXPTIME-complete. For practical purposes, we leverage from recent advances in automata theory, such as inclusion checking for non-deterministic TA avoiding explicit determinization. We implemented our method and present promising preliminary experimental results.

1 Introduction

Separation Logic (SL) [17] is a logical framework for describing recursive mutable data structures. The attractiveness of SL as a specification formalism comes from the possibility of writing higher-order *inductive definitions* that are natural for describing the most common recursive data structures, such as singly- or doubly-linked lists (SLLs/DLLs), trees, hash maps (lists of lists), and more complex variations thereof, such as nested and overlaid structures (e.g. lists with head and tail pointers, skip-lists, trees with linked leaves, etc.). In addition to being an appealing specification tool, SL is particularly suited for compositional reasoning about programs. Indeed, the principle of *local reasoning* allows one to verify different elements (functions, threads) of a program, operating on disjoint parts of the memory, and to combine the results a-posteriori, into succinct verification conditions.

However, the expressive power of SL comes at the price of undecidability [6]. To avoid this problem, most SL dialects used by various tools (e.g. SPACE INVADER [3], PREDATOR [10], or INFER [7]) use hard-coded predicates, describing SLLs and DLLs, for which entailments are, in general, tractable [9]. For graph structures of bounded tree width, a general decidability result was presented in [13]. Entailment in this fragment is EXPTIME-hard as proved in [1].

In this paper, we present a novel decision procedure for a restriction of the decidable fragment of [13] that describes recursive structures in which all edges are *local* with respect to a spanning tree. Examples of such structures include SLLs, DLLs, trees and trees with parent pointers, etc. For structures outside of this class (e.g. skip-lists or trees with linked leaves), our procedure is sound (the answer is positive if the entailment holds), but not complete (the answer might be negative and the entailment could still hold). In terms of program verification, such a lack of completeness in the entailment prover can cause non-termination of the fixpoint, but will not provide unsound results.

The method described in the paper belongs to the class of *automata-theoretic* decision techniques: we translate an entailment problem $\varphi \models \psi$ into a language inclusion

problem $\mathcal{L}(A_\varphi) \subseteq \mathcal{L}(A_\psi)$ for tree automata (TA) A_φ and A_ψ that (roughly speaking) encode the sets of models of φ and ψ, respectively. Yet, a naïve translation of the inductive definitions of SL into TA encounters a *polymorphic representation* problem: the same structure can be defined in several different ways, and TA simply mirroring the definition will not report the entailment. For example, DLLs with selectors `next` and `prev` for the next and previous nodes, respectively, can be described by a forward unfolding of the inductive definition $\mathtt{DLL}(\mathit{head}, \mathit{prev}, \mathit{tail}, \mathit{next}) \equiv \exists x.\ \mathit{head} \mapsto (x, \mathit{prev}) * \mathtt{DLL}(x, \mathit{head}, \mathit{tail}, \mathit{next})$ as well as by a backward unfolding of the definition $\mathtt{DLL}_{\mathit{rev}}(\mathit{head}, \mathit{prev}, \mathit{tail}, \mathit{next}) \equiv \exists x.\ \mathit{tail} \mapsto (\mathit{next}, x) * \mathtt{DLL}_{\mathit{rev}}(\mathit{head}, \mathit{prev}, x, \mathit{tail})$. Also, one can define a DLL starting with a node in the middle and unfolding backward to the left of this node and forward to the right: $\mathtt{DLL}_{\mathit{mid}}(\mathit{head}, \mathit{prev}, \mathit{tail}, \mathit{next}) \equiv \exists x, y, z\ .\ x \mapsto (y, z) * \mathtt{DLL}(y, x, \mathit{tail}, \mathit{next}) * \mathtt{DLL}_{\mathit{rev}}(\mathit{head}, \mathit{prev}, z, x)$. The entailments $\mathtt{DLL}(\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}) \models \mathtt{DLL}_{\mathit{rev}}(\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d})$ and $\mathtt{DLL}_{\mathit{mid}}(\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}) \models \mathtt{DLL}(\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d})$ hold, but a naïve structural translation to TA might not detect this fact. To bridge this gap, we define a closure operation on TA, called *canonical rotation*, which automatically adds all possible representations of a given inductive definition, encoded as a tree automaton.

Our reduction from SL to TA provides tight complexity bounds showing that entailment in our fragment is EXPTIME-complete. Moreover, from a practical point view, we implemented our method using the VATA [14] tree automata library, which leverages from recent advances in non-deterministic language inclusion for TA [5], and obtained quite encouraging experimental results.

Related work. Given the large body of literature on decidable logics for describing mutable data structures, we need to restrict this section to the related work that focuses on SL [17]. The first (proof-theoretic) decidability result for SL on a restricted fragment defining only SLLs was reported in [4], which describe a co-NP algorithm. The full basic SL without recursive definitions, but with the magic wand operator was found to be undecidable when interpreted *in any memory model* [6]. A PTIME entailment procedure for SL with list predicates is given in [9]. Their method was extended to reason about nested and overlaid lists in [11]. More recently, entailments in an important SL fragment with hardcoded SLL/DLL predicates were reduced to Satisfiability Modulo Theories (SMT) problems, leveraging from recent advances in SMT technology [16].

Closer to our work on SL with user-provided *inductive definitions* is the fragment used in the tool SLEEK, which implements a semi-algorithmic entailment check, based on unfoldings and unifications [15]. Our previous work [13] gives a general decidability result for SL with inductive definitions interpreted over graph-like structures, under several necessary restrictions. The work [1] provides a rather complete picture of complexity for the entailment in various SL fragments with inductive definitions, including EXPTIME-hardness of the decidable fragment of [13], but provides no upper bound. The EXPTIME-completness result in this paper provides an upper bound for the fragment of *local* definitions, and strengthens the EXPTIME-hard lower bound as well.

2 Definitions

The set of natural numbers is denoted by $\mathbb{N}$. For a finite set S, we denote by $\|S\|$ its cardinality. If $\mathbf{x} = \langle x_1, \ldots, x_n \rangle$ and $\mathbf{y} = \langle y_1, \ldots, y_m \rangle$ are tuples, $\mathbf{x} \cdot \mathbf{y} = \langle x_1, \ldots, x_n, y_1, \ldots, y_m \rangle$ denotes their concatenation, and $(\mathbf{x})_i = x_i$ denotes the i-th element of $\mathbf{x}$. If x is an element, we denote by $\mathbf{x}_{\neg x}$ the result of projecting out all occurrences of x from $\mathbf{x}$, and if S

is a set, we denote by $\mathbf{x}_{\neg S}$ the result of projecting out all occurrences of some element $s \in S$ from $\mathbf{x}$. The intersection $\mathbf{x} \cap S$ denotes the tuple in which only elements of S are maintained. The union (intersection) $\mathbf{x} \cup \mathbf{y}$ ($\mathbf{x} \cap \mathbf{y}$) denotes the set of elements from either (both) $\mathbf{x}$ and $\mathbf{y}$. For two tuples of variables $\mathbf{x} = \langle x_1, \ldots, x_k \rangle$ and $\mathbf{y} = \langle y_1, \ldots, y_k \rangle$, of equal length, we write $\mathbf{x} = \mathbf{y}$ for $\bigwedge_{i=1}^{k} x_i = y_i$. Moreover, for a single variable z, we write $\mathbf{x} = z$ for $\bigwedge_{i=1}^{k} x_i = z$.

For a partial function $f : A \rightharpoonup B$, and $\bot \notin B$, we denote $f(x) = \bot$ the fact that f is undefined at some point $x \in A$. The domain of f is denoted $dom(f) = \{x \in A \mid f(x) \neq \bot\}$, and the image of f is denoted as $img(f) = \{y \in B \mid \exists x \in A \,.\, f(x) = y\}$. By $f : A \rightharpoonup_{fin} B$, we denote any partial function whose domain is finite. Given two partial functions f, g defined on disjoint domains, we denote by $f \oplus g$ their union.

2.1 Stores, Heaps, and States

We consider $Var = \{x, y, z, \ldots\}$ to be a countably infinite set of *variables* and $\mathbf{nil} \in Var$ be a designated variable. We assume a total lexicographical ordering on the set of variables Var, and for any set $S \subseteq Var$, we denote by $minlex(S)$ the unique minimal element with respect to this ordering. Let Loc be a countably infinite set of locations, $null \in Loc$ be a designated location, and $Sel = \{1, \ldots, \mathcal{S}\}$, for some given $\mathcal{S} > 0$, be a finite set of natural numbers, called *selectors* in the following.

Definition 1. *A* state *is a pair $\langle s, h \rangle$ where $s : Var \rightharpoonup Loc$ is a partial function mapping pointer variables into locations such that $s(\mathbf{nil}) = null$, and $h : Loc \rightharpoonup_{fin} Sel \rightharpoonup_{fin} Loc$ is a finite partial function such that (i) $null \notin dom(h)$ and (ii) for all $\ell \in dom(h)$ there exists $k \in Sel$ such that $(h(\ell))(k) \neq \bot$.*

Given a state $S = \langle s, h \rangle$, s is called the *store* and h the *heap*. For any $k \in Sel$, we write $\ell \xrightarrow{k}_S \ell'$ for $(h(\ell))(\sigma) = \ell'$. We call a triple $\ell \xrightarrow{k}_S \ell'$ an *edge* of S. Sometimes we omit the subscript when it is obvious from the context. Let $Img(h) = \bigcup_{\ell \in Loc} img(h(\ell))$ be the set of locations which are destinations of some edge in h. A location $\ell \in Loc$ is said to be *allocated* in $\langle s, h \rangle$ if $\ell \in dom(h)$ (i.e. it is the source of an edge), and *dangling* in $\langle s, h \rangle$ if $\ell \in [img(s) \cup Img(h)] \setminus dom(h)$, i.e. it is referenced by a store variable, or reachable from an allocated location in the heap, but it is not allocated in the heap itself. The set $loc(S) = img(s) \cup dom(h) \cup Img(h)$ is the set of all locations either allocated or referenced in the state S.

For any two states $S_1 = \langle s_1, h_1 \rangle$ and $S_2 = \langle s_2, h_2 \rangle$, such that (i) s_1 and s_2 agree on the evaluation of common variables ($\forall x \in dom(s_1) \cap dom(s_2) \,.\, s_1(x) = s_2(x)$), and (ii) h_1 and h_2 have disjoint domains ($dom(h_1) \cap dom(h_2) = \emptyset$), we denote by $S_1 \uplus S_2 = \langle s_1 \cup s_2, h_1 \oplus h_2 \rangle$ the *disjoint union* of S_1 and S_2. The disjoint union is undefined if one of the above conditions does not hold.

For a state $S = \langle s, h \rangle$ and a location $\ell \in dom(h)$, the *neighbourhood of ℓ in S* is a state denoted as $S_{\langle \ell \rangle} = \langle s_\ell, h_\ell \rangle$, where:

- $h_\ell = \{\langle \ell, \lambda k \,.\, \text{if } \ell \xrightarrow{k}_S \ell' \text{ then } \ell' \text{ else } \bot \rangle\}$
- $s_\ell(x) =$ if $s(x) \in dom(h_\ell) \cup img(h_\ell)$ then $s(x)$ else $\bot$

Intuitively, the neighbourhood of an allocated location ℓ is the state in which only ℓ is allocated and all other locations ℓ' for which there is an edge $\ell \xrightarrow{k}_S \ell'$ are dangling.

2.2 Trees

Let Σ be a countable alphabet, and $\mathbb{N}^*$ be the set of sequences of natural numbers. Let $\varepsilon \in \mathbb{N}^*$ denote the empty sequence, and $p.q$ denote the concatenation of two sequences $p,q \in \mathbb{N}^*$. We say that p is a *prefix* of q if $q = p.q'$, for some $q' \in \mathbb{N}^*$.

The total order on $\mathbb{N}$ extends to a total lexicographical order on $\mathbb{N}^*$, and for any set $S \subseteq \mathbb{N}^*$, we denote by $minlex(S)$[1] the lexicographically minimal element of S.

A *tree* t over Σ is a finite partial function $t : \mathbb{N}^* \rightharpoonup_{fin} \Sigma$ such that $dom(t)$ is a finite prefix-closed subset of $\mathbb{N}^*$ and, for each $p \in dom(t)$ and $i \in \mathbb{N}$, we have $t(p.i) \neq \bot$ only if $t(p.j) \neq \bot$, for all $0 \leq j < i$. The sequences $p \in dom(t)$ are called *positions* in the following. Given two positions $p,q \in dom(t)$, we say that q is the i-th successor (child) of p if $q = p.i$, for $i \in \mathbb{N}$.

We denote by $\mathcal{D}(t) = \{-1,0,\ldots,N\}$ the *direction alphabet* of t, where $N = \max\{i \in \mathbb{N} \mid \exists p \in \mathbb{N}^* \,.\, p.i \in dom(t)\}$, and we let $\mathcal{D}_+(t) = \mathcal{D}(t) \setminus \{-1\}$. By convention, we have $(p.i).(-1) = p$, for all $p \in \mathbb{N}^*$ and $i \in \mathcal{D}_+(t)$. Given a tree t and a position $p \in dom(t)$, we define the *arity* of the position p as $\#_t(p) = \max\{d \in \mathcal{D}_+(t) \mid p.d \in dom(t)\} + 1$, and the *subtree of t rooted at p* as $t_{|_p}(q) = t(p.q)$, for all $q \in \mathbb{N}^*$.

A *path* in t, from p_1 to p_k, is a sequence $p_1, p_2, \ldots, p_k \in dom(t)$ of pairwise distinct positions, such that, for all $1 \leq i < k$, there exist $d_i \in \mathcal{D}(t)$ such that $p_{i+1} = p_i.d_i$. Notice that a path in the tree can also link sibling nodes, not just ancestors to their descendants or vice versa. However, a path may not visit the same tree position twice.

Given a position $p \in dom(t)$, we denote by $t_{|_p}$ the subtree of t starting at position p, i.e. $t_{|_p}(q) = t(p.q)$, for all $q \in \mathbb{N}^*$.

Definition 2. *Given two trees $t_1, t_2 : \mathbb{N}^* \rightharpoonup_{fin} \Sigma$, we say that t_2 is a* rotation *of t_1 denoted by $t_1 \sim_r t_2$ if and only if $r : dom(t_1) \to dom(t_2)$ is a bijective function such that: $\forall p \in dom(t_1) \forall d \in \mathcal{D}_+(t_1) \,:\, p.d \in dom(t_1) \Rightarrow \exists e \in \mathcal{D}(t_2) \,.\, r(p.d) = r(p).e$. We write $t_1 \sim t_2$ if there exists a function $r : dom(t_1) \to dom(t_2)$ such that $t_1 \sim_r t_2$.*

An example of a rotation r of a tree t_1 to a tree t_2 such that $r(\varepsilon) = 2$, $r(0) = \varepsilon$, $r(1) = 20$, $r(00) = 0$, and $r(01) = 1$ is shown in Fig. 1. Note that, e.g., for $p = \varepsilon \in dom(t_1)$ and $d = 0 \in \mathcal{D}_+(t_1)$ where $p.d = \varepsilon.0 \in dom(t_1)$, we get $e = -1 \in \mathcal{D}(t_2)$ and $r(\varepsilon.0) = 2.(-1) = \varepsilon$.

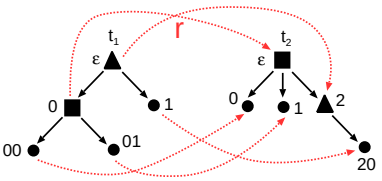

Fig. 1. An example of a rotation.

Proposition 1. *The relation $\sim$ is an equivalence relation.*

Proof. The relation $\sim$ is clearly reflexive as one can choose r as the identity function. To prove that $\sim$ is transitive, let $t_1, t_2, t_3 : \mathbb{N}^* \rightharpoonup_{fin} \Sigma$ be trees such that $t_1 \sim_{r_1} t_2$ and $t_2 \sim_{r_2} t_3$, and let $r_1 : dom(t_1) \to dom(t_2)$ and $r_2 : dom(t_2) \to dom(t_3)$ be the bijective functions from Def. 2, respectively. Then, for all $p \in dom(t_1)$ and $d \in \mathcal{D}_+(t_1)$ such that $p.d \in dom(t_1)$, there exists $e_1 \in \mathcal{D}(t_2)$ such that $r_1(p.d) = r_1(p).e_1$. We distinguish two cases:

1. If $e_1 \in \mathcal{D}_+(t_2)$, then there exists $e_2 \in \mathcal{D}(t_3)$ such that $(r_2 \circ r_1)(p.d) = r_2(r_1(p).e_1) = (r_2 \circ r_1)(p).e_2$.

[1] The distinction with the lexicographical order on *Var* is clear from the type of elements in S.

2. Otherwise, if $e_1 = -1$, then $r_1(p) = r_1(p.d).e'_1$, for some $e'_1 \in \mathcal{D}_+(t_2)$. Then there exists $e_2 \in \mathcal{D}(t_3)$ such that $(r_2 \circ r_1)(p) = r_2(r_1(p.d).e'_1) = (r_2 \circ r_1)(p.d).e_2$, and consequently, $(r_2 \circ r_1)(p.d) = (r_2 \circ r_1)(p).(-1)$.

To show that $\sim$ is symmetric, let $t_1 \sim_r t_2$ be two trees, and let $p \in dom(t_2)$ and $d \in \mathcal{D}_+(t_2)$ such that $p.d \in dom(t_2)$. We prove that there exists $e \in \mathcal{D}(t_1)$ such that $r^{-1}(p.d) = r^{-1}(p).e$, where r is the bijective function from Def. 2. By contradiction, suppose that there is no such direction, and since $r^{-1}(p), r^{-1}(p.d) \in dom(t_1)$, there exists a path $r^{-1}(p) = p_1, \ldots, p_i, \ldots, p_k = r^{-1}(p.d)$ of length $k > 2$ in t_1 such that:

- $p_{j+1} = p_j.(-1)$, for all $1 \leq j < i$,
- $p_{j+1} = p_j.d_j$, for all $i \leq j < k$ and $d_i, \ldots, d_{k-1} \in \mathcal{D}_+(t_1)$

where $p_1, \ldots, p_k$ are distinct positions in $dom(t_1)$. Then $r(p_1), \ldots, r(p_k)$ are distinct positions in $dom(t_2)$ such that:

- $r(p_j) = r(p_{j+1}).e_j$, for all $1 \leq j < i$ and some $e_j \in \mathcal{D}(t_2)$,
- $r(p_{j+1}) = r(p_j).e_j$, for all $i \leq j < k$ and some $e_j \in \mathcal{D}(t_2)$.

So there exists a path $p = r(p_1), \ldots, r(p_k) = p.d$ of length $k > 2$ in t_2, which contradicts with the fact that t_2 is a tree. □

A rotation function r between two trees $t \sim_r u$ is said to *revert* two positions $p, q \in dom(t)$, if q is a prefix of p and $r(p)$ is a prefix of $r(q)$. We show next that the only reversions in $dom(t)$ due to r appear on the (unique) path from the root ε to $r^{-1}(\varepsilon)$.

Lemma 1. *Let $t, u : \mathbb{N}^* \rightharpoonup_{fin} \Sigma$ be two trees, $r : dom(t) \to dom(u)$ be a bijective function such that $t \sim_r u$, and $p \in dom(t)$ be a position such that $r(p) = \varepsilon$ is the root of u. Then, for all $q \in dom(t)$ and all $0 \leq d < \#_t(q)$, $r(q.d) = r(q).(-1)$ iff $q.d$ is a prefix of p.*

Proof. "⇒" Let $q \in dom(t)$ be an arbitrary position such that $r(q.d) = r(q).(-1)$, for some $0 \leq d < \#_t(q)$. Suppose, by contradiction, that $q.d$ is not a prefix of p. There are two cases:

1. p is a strict prefix of $q.d$, i.e. there exists a sequence $p = p_0, \ldots, p_k = q$, for some $k > 0$, such that, for all $0 \leq i < k$, there exists $0 \leq j_i < \#_t(p_i)$ such that $p_{i+1} = p_i.j_i$. Then, there exists $\ell \in \mathcal{D}(u)$ such that $r(q) = r(p_{k-1}.j_{k-1}) = r(p_{k-1}).\ell$. If $\ell \geq 0$, we have $r(p_{k-1}) = r(q).(-1) = r(q.d) = r(p_{k-1}.j_{k-1}.d)$. Since both $j_{k-1}, d \geq 0$, p_{k-1} and $p_{k-1}.j_{k-1}.d$ are distinct nodes from $dom(t)$, and we reach a contradiction with the fact that r is bijective. Hence, $\ell = -1$ is the only possibility. Applying the same argument inductively on $p_0, \ldots, p_k$, we find that $r(p)$ is either equal or is a descendant of $r(q)$, which contradicts with the fact that $r(p) = \varepsilon$ is the root of u.
2. p is not a prefix of $q.d$, and there exists a common prefix p_0 of both p and $q.d$. Let p_0 be the maximal such prefix. Since p_0 is a prefix of $q.d$, there exists a sequence $p_0, \ldots, p_k = q$, for some $k > 0$, such that, for all $0 \leq i < k$ there exists $0 \leq j_i < \#_t(p_i)$ such that $p_{i+1} = p_i.j_i$. By the argument of the previous case, we have that $r(q) = r(p_k), r(p_{k-1}), \ldots, r(p_0)$ is a strictly descending path in u, i.e. $r(p_i)$ is a child of $r(p_{i+1})$, for all $0 \leq i < k$. Since p_0 is a prefix of p, there exists another non-trivial sequence $p_0 = p'_0, \ldots, p'_m = p$, for some $m > 0$, such that,

for all $0 \leq i < m$ there exists $0 \leq j'_i < \#_t(p'_i)$ and $p'_{i+1} = p'_i.j'_i$. Moreover, since p_0 is the maximal prefix of p and $q.d$, we have $\{p_1, \ldots, p_k\} \cap \{p'_1, \ldots, p'_m\} = \emptyset$. Then $r(p'_1) = r(p_0.j'_0) = r(p_0).\ell$, for some $\ell \in \mathcal{D}(u)$. If $\ell = -1$, then necessarily $r(p'_1) = r(p_{k-1})$, which contradicts with the fact that r is a bijection, since $p'_1 \neq p_{k-1}$. Then the only possibility is $\ell \geq 0$. By applying the same argument inductively on $p'_0, \ldots, p'_m = p$, we obtain that $r(p)$ is a strict descendant of $r(p_0)$, which contradicts with $r(p) = \varepsilon$ being the root of u.

Since both cases above lead to contradictions, the only possibility is that $q.d$ is a prefix of p.

"$\Leftarrow$" If $q.d$ is a prefix of p, there exists a non-trivial sequence $q = q_0, q.d = q_1, \ldots, q_k = p$, for some $k > 0$, such that, for all $0 \leq i < k$, there exists $0 \leq j_i < \#_t(q_i)$ such that $q_{i+1} = q_i.j_i$. Then, $\varepsilon = r(p) = r(q_k) = r(q_{k-1}.j_{k-1}) = r(q_{k-1}).(-1)$. Reasoning inductively, we obtain that, for all $0 < i \leq k$, $r(q_i) = r(q_{i-1}).(-1)$, hence $r(q.d) = r(q_1) = r(q_0).(-1) = r(q).(-1)$. □

A rotation function r between two trees $t \sim_r u$ is said to *revert* two positions $p, q \in dom(t)$, if q is a prefix of p and $r(p)$ is a prefix of $r(q)$. One can see that the only reversions in $dom(t)$ due to r appear on the (unique) path from the root ε to $r^{-1}(\varepsilon)$.

Definition 3. *Given a state $S = \langle s, h \rangle$, a* spanning tree *of S is a bijective tree $t : \mathbb{N}^* \to dom(h)$ such that $\forall p \in dom(t) \forall d \in \mathcal{D}_+(t)$. $p.d \in dom(t) \Rightarrow \exists k \in Sel$. $t(p) \xrightarrow{k}_S t(p.d)$. An edge $\ell \xrightarrow{k}_S \ell'$ is said to be* local *with respect to a spanning tree t iff there exist $p \in dom(t)$ and $d \in \mathcal{D}(t) \cup \{\varepsilon\}$ such that $t(p) = \ell$ and $t(p.d) = \ell'$. Moreover, t is a* local spanning tree *of S if t is a spanning tree of S, and S has only local edges wrt. t.*

An example of a doubly-linked list and two of its spanning trees is shown in Fig. 2. Note that both the `next` and `prev` edges are local in both cases (they get mapped either to the 0 or -1 direction) and so both of the spanning trees are local.

Notice that a spanning tree covers exactly the set of allocated locations in a state. Moreover, if a state S has a spanning tree t, then every location $\ell \in loc(S)$, not necessarily allocated, is reachable from $t(\varepsilon)$ by a chain of edges. In other words, the state does not contain garbage nodes.

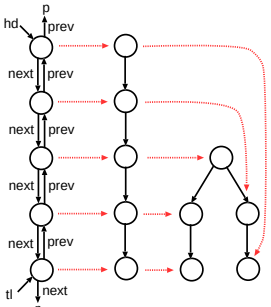

Fig. 2. Spanning trees.

Lemma 2. *Let $S = \langle s, h \rangle$ be a state and t its spanning tree. If all edges of S are local with respect to t, then for each spanning tree t' of S we have $t \sim_{t'^{-1} \circ t} t'$, and, moreover, all edges of S are local with respect to t' as well.*

Proof. Let $p \in dom(t)$ be an arbitrary position and $d \in \mathcal{D}_+(t)$ be an arbitrary non-negative direction such that $p.d \in dom(t)$. If t is a spanning tree, then t is bijective, and there exists a selector $s \in Sel$ such that $t(p) \xrightarrow{s} t(p.d)$ is an edge in S. Since t' is a spanning tree, it is also bijective, hence there exist $p', p'' \in dom(t')$ such that $t'(p') = t(p)$ and $t'(p'') = t(p.d)$.

By contradiction, suppose that $p'' \neq p'.e$ for all $e \in \mathcal{D}(t')$. Since $p', p'' \in dom(t')$, there exists a path $p' = p_1, \ldots, p_k$ such that $k > 2$, and, for all $1 \leq i < k$, we have $p_{i+1} = p_i.e_i$, for some $e_1, \ldots, e_{k-1} \in \mathcal{D}(t')$. Because t' is a spanning tree of S, there exist selectors $s_1, \ldots, s_{k-1} \in Sel$ such that one of the following holds, for all $1 \leq i < k$:

- $t'(p_i) \xrightarrow{s_i} t'(p_{i+1})$ or
- $t'(p_{i+1}) \xrightarrow{s_i} t'(p_i)$.

Since all the above edges of S are local wrt. t, there exists a path $p = (t^{-1} \circ t')(p_1), \ldots, (t^{-1} \circ t')(p_k) = p.d$ of pairwise distinct positions in t, for $k > 2$, which contradicts with the fact that t is a tree. In conclusion, for all $p \in dom(t)$ and $d \in \mathcal{D}_+(t)$, such that $p.d \in dom(t)$ there exists $e \in \mathcal{D}(t')$ such that $(t'^{-1} \circ t)(p.d) = (t'^{-1} \circ t)(p).e$, and since $t'^{-1} \circ t$ is a bijective mapping, we have $t \sim_{t'^{-1} \circ t} t'$.

Finally, we are left with proving that all edges are local with respect to t'. Let $\ell \xrightarrow{s} \ell'$ be an arbitrary edge of S, for some $s \in Sel$. The edge is local with respect to t, thus there exists $p \in dom(t)$ and $d \in \mathcal{D}(t) \cup \{\varepsilon\}$ such that $t(p) = \ell$ and $t(p.d) = \ell'$. We distinguish three cases:

1. If $d = \varepsilon$, then $p = p.d$, and trivially $\ell = t'((t'^{-1} \circ t)(p)) = t'((t'^{-1} \circ t)(p.d)) = t'((t'^{-1} \circ t)(p).d) = \ell'$.
2. Otherwise, if $d \in \mathcal{D}_+(t)$, taking into account that $t \sim_{t'^{-1} \circ t} t'$, by the first part of this lemma, we have $(t'^{-1} \circ t)(p.d) = (t'^{-1} \circ t)(p).e$, for some $e \in \mathcal{D}(t')$. But $t'((t'^{-1} \circ t)(p)) = t(p) = \ell$ and $t'((t'^{-1} \circ t)(p.d)) = t(p.d) = \ell'$, so $\ell \xrightarrow{s} \ell'$ is local wrt. t'.
3. Otherwise, if $d = -1$, there exists $q \in dom(t)$ and $d' \in \mathcal{D}_+(t)$ such that $p = q.d'$. Since $t \sim_{t'^{-1} \circ t} t'$, by the first part of this lemma, we have $(t'^{-1} \circ t)(q.d') = (t'^{-1} \circ t)(q).e$, for some $e \in \mathcal{D}(t')$. But $t'(t'^{-1} \circ t)(q.d') = t(p) = \ell$ and $t'((t'^{-1} \circ t)(q)) = t(q) = t(p.d) = \ell'$, so $\ell \xrightarrow{s} \ell'$ is local with respect to t'.

□

As a consequence, if all edges of a state are local with respect to some spanning tree, they are also local with respect to any other spanning tree, hence we will simply say that they are local. A state is said to be *local* if it has only local edges.

2.3 Separation Logic

The syntax of *basic formulae* of Separation Logic (SL) is given below:

$$\begin{array}{l} \alpha \ \in \ Var \setminus \{\mathbf{nil}\};\ x \ \in \ Var; \\ \Pi ::= \alpha = x \mid \Pi_1 \wedge \Pi_2 \\ \Sigma \ ::= \mathbf{emp} \mid \alpha \mapsto (x_1, \ldots, x_n) \mid \Sigma_1 * \Sigma_2 \text{ , for some } n > 0 \\ \varphi \ ::= \Sigma \wedge \Pi \mid \exists x \,.\, \varphi_1 \end{array}$$

A formula of the form $\bigwedge_{i=1}^{n} \alpha_i = x_i$ defined by the Π nonterminal in the syntax above is said to be *pure*. A formula of the form $\bigstar_{i=1}^{k} \alpha_i \mapsto (x_{i,1}, \ldots, x_{i,n})$ defined by the Σ nonterminal in the syntax above is said to be *spatial*. The atomic proposition **emp** denotes the empty spatial conjunction. A variable x is said to be *free* in φ if it does not occur under the scope of any existential quantifier. We denote by $FV(\varphi)$ the set of free variables, and by $AP(\varphi)$ the set of atomic propositions of φ.

In the following, we shall use two equality relations. The *syntactic equality*, denoted $\alpha \equiv \beta$, means that α and β are the same syntactic object (formula, variable, tuples of variables, etc.). On the other hand, by writing $x =_\Pi y$, for two variables $x, y \in Var$ and a pure formula Π, we mean that the equality of the values of x and y is implied by Π.

A *substitution* is an injective partial function $\sigma : Var \rightharpoonup_{fin} Var$. Given a basic formula φ and a substitution σ, we denote by $\varphi[\sigma]$ the result of simultaneously replacing each variable (not necessarily free) x that occurs in φ, by $\sigma(x)$. For instance, if $\sigma(x) = y$, $\sigma(y) = z$ and $\sigma(z) = t$, then $(\exists x, y \,.\, x \mapsto (y, z) \wedge z = x)[\sigma] \equiv \exists y, z \,.\, y \mapsto (z, t) \ \wedge \ t = y$.

The semantics of a basic formula φ is given by the relation $S \models \varphi$, where $S = \langle s, h \rangle$ is a state such that $FV(\varphi) \subseteq dom(s)$, and φ is a basic SL formula. The definition of $\models$ is by induction on the structure of φ:

$$
\begin{array}{lll}
S \models \mathbf{emp} & \iff & dom(h) = \emptyset \\
S \models \alpha \mapsto (x_1, \ldots, x_n) & \iff & s = \{(\alpha, \ell_0), (x_1, \ell_1), \ldots, (x_n, \ell_n)\} \text{ and} \\
 & & h = \{\langle \ell_0, \lambda i \,.\, \text{if } 1 \leq i \leq n \text{ then } \ell_i \text{ else } \bot \rangle\} \\
 & & \text{for some } \ell_0, \ell_1, \ldots, \ell_n \in Loc \\
S \models \varphi_1 * \varphi_2 & \iff & S_1 \models \varphi_1 \text{ and } S_2 \models \varphi_2 \text{ where } S_1 \uplus S_2 = S \\
S \models \exists x \,.\, \varphi & \iff & \langle s[x \leftarrow \ell], h \rangle \models \varphi \text{ for some } \ell \in Loc
\end{array}
$$

The semantics of $=$ and $\wedge$ is classical in first order logic. Note that we adopt here the *strict semantics*, in which a points-to relation $\alpha \mapsto (x_1, \ldots, x_n)$ holds in a state consisting of a single cell pointed to by α, with exactly n outgoing edges $s(\alpha) \xrightarrow{k}_S s(x_k)$, $1 \leq k \leq n$, towards either the single allocated location (if $s(x_k) = s(\alpha)$), or dangling locations (if $s(x_k) \neq s(\alpha)$). The empty heap is specified by **emp**.

A variable $x \in FV(\Sigma)$ is said to be *allocated* in a basic spatial formula Σ if it occurs on the right hand side of a points-to atomic proposition $x \mapsto (y_0, \ldots, y_{k-1})$ of Σ. If Σ is satisfiable, then clearly each free variable $x \in FV(\Sigma)$ is allocated at most once. For a basic quantifier-free SL formula $\varphi \equiv \Sigma \wedge \Pi$ and two variables $x, y \in FV(\varphi)$, we say that y is φ-*reachable* from x in iff there exists a sequence $x =_\Pi \alpha_0, \ldots, \alpha_m =_\Pi y$, for some $m \geq 0$, such that, for each $0 \leq i < m$, $\alpha_i \mapsto (\beta_{i,1}, \ldots, \beta_{i,p_i})$ is a points-to proposition in Σ, and $\beta_{i,s} =_\Pi \alpha_{i+1}$, for some $1 \leq s \leq p_i$. A variable $x \in FV(\Sigma)$ is said to be a *root* of Σ if every variable $y \in FV(\Sigma)$ is reachable from x.

Remark. Notice that there is no explicit disequality between variables in the basic fragment of SL. This is in part justified by the fact that disequality can be partially defined using the following implication: $\alpha \mapsto (x_1, \ldots, x_n) * \beta \mapsto (y_1, \ldots, y_m) \Rightarrow \alpha \neq \beta$. This implication enforces disequality only between allocated variables. However, in practice, most SL specifications do not impose disequality constraints on dangling variables.

2.4 Inductive Definitions

A system of *inductive definitions*[2] (inductive system) $\mathcal{P}$ is a set of rules of the form:

$$
\begin{array}{rcl}
P_1(x_{1,1}, \ldots, x_{1,n_1}) & \equiv & |_{j=1}^{m_1} R_{1,j}(x_{1,1}, \ldots, x_{1,n_1}) \\
 & \ldots & \\
P_k(x_{k,1}, \ldots, x_{k,n_k}) & \equiv & |_{j=1}^{m_k} R_{k,j}(x_{k,1}, \ldots, x_{k,n_k})
\end{array}
\tag{1}
$$

where $\{P_1, \ldots, P_k\}$ is a set of *predicates*, $x_{i,1}, \ldots, x_{i,n_i}$ are called *formal parameters*, and the formulae $R_{i,j}$ are called the *rules* of P_i. Concretely, each rule is of the form:

$$R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \Sigma * P_{i_1}(\mathbf{y}_1) * \ldots * P_{i_m}(\mathbf{y}_m) \ \wedge \ \Pi$$

[2] The name *inductive* suggests that every structure described by the system is finite.

where $\mathbf{x} \cap \mathbf{z} = \emptyset$, and all of the following hold:

1. $\mathbf{head}(R_{i,j}) \stackrel{def}{\equiv} \Sigma$ is a non-empty spatial formula (i.e. $\Sigma \not\equiv \mathbf{emp}$), and $FV(\Sigma) \subseteq \mathbf{x} \cup \mathbf{z}$,
2. $\mathbf{tail}(R_{i,j}) \stackrel{def}{\equiv} \langle P_{i_1}(\mathbf{y}_1), \ldots, P_{i_m}(\mathbf{y}_m) \rangle$ is an ordered sequence of *predicate occurrences*, where $\mathbf{y}_1 \cup \ldots \cup \mathbf{y}_m \subseteq \mathbf{x} \cup \mathbf{z}$,
3. Π is a pure formula, and $FV(\Pi) \subseteq \mathbf{z} \cup \mathbf{x}$. In the following, we restrict the pure part of each rule such that, for all formal parameters $\beta \in \mathbf{x}$, we allow only equalities of the form $\alpha =_\Pi \beta$, where α is allocated in Σ. This restriction is of technical nature (see Section 4.1). It is possible to lift it, but only at the expense of an exponential blowup in the size of the resulting tree automaton.
4. For all $1 \leq r, s \leq m$, if $x_{i,k} \in \mathbf{y}_r$, $x_{i,l} \in \mathbf{y}_s$, and $x_{i,k} =_\Pi x_{i,l}$, for some $1 \leq k, l \leq n_i$, then $r = s$. In other words, a formal parameter of a rule cannot be passed to two or more subsequent occurrences of predicates in that rule. This technical restriction can be lifted at the cost of introducing expensive tests for double allocation as was done in the translation of the inductive definitions to Monadic Second-Order Logic on graphs, reported in [13].

The *size* of a rule R is denoted by $|R|$ and is defined inductively as follows:

$$|\alpha = x| = 1 \qquad |\mathbf{emp}| = 1 \qquad |\alpha \mapsto (x_1, \ldots, x_n)| = n+1$$
$$|\varphi \bullet \psi| = |\varphi| + |\psi| \quad |\exists x \,.\, \varphi| = |\varphi| + 1 \qquad |P(x_1, \ldots, x_n)| = n$$

where $\alpha \in \mathit{Var} \setminus \{\mathbf{nil}\}$, $x, x_1, \ldots, x_n \in \mathit{Var}$, and $\bullet \in \{*, \wedge\}$. The size of an inductive system $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ is defined as $|\mathcal{P}| = \sum_{i=1}^{n} \sum_{j=1}^{m_i} |R_{i,j}|$.

Example 1. To illustrate the use inductive definitions (with the above restrictions), we first show how to define a predicate $\mathtt{DLL}(hd, p, tl, n)$ describing doubly-linked lists of length at least one. As depicted on the left of Fig. 2, the formal parameter *hd* points to the first allocated node of such a list, *p* to the node pointed by the *prev* selector of *hd*, *tl* to the last node of the list (possibly equal to *hd*), and *n* to the node pointed by the *next* selector from *tl*. This predicate can be defined as follows:

$$\mathtt{DLL}(hd, p, tl, n) \equiv hd \mapsto (n, p) \;\wedge\; hd = tl \;\mid\; \exists x.\; hd \mapsto (x, p) * \mathtt{DLL}(x, hd, tl, n)$$

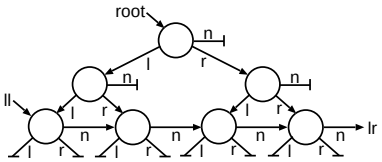

Fig. 3. A tree with linked leaves.

Example 2. Another example is the predicate $\mathtt{TLL}(root, ll, lr)$ describing binary trees with linked leaves, whose root is pointed by the *root* formal parameter, the left-most leaf is pointed to by *ll*, and the right-most leaf points to *lr* (cf. Fig. 3):

$$\mathtt{TLL}(root, ll, lr) \equiv root \mapsto (\mathbf{nil}, \mathbf{nil}, lr) \;\wedge\; root = ll$$
$$\mid\; \exists x, y, z.\; root \mapsto (x, y, \mathbf{nil}) * \mathtt{TLL}(x, ll, z) * \mathtt{TLL}(y, z, lr)$$

Remark. The first restriction above ($\mathbf{head}(R_{i,j}) \not\equiv \mathbf{emp}$) currently does not allow defining data structures of zero size (such as empty lists, trees, etc.) and, more importantly, certain types of structure concatenation. For instance, considering the DLL predicate from the above example, the following predicate cannot be defined in our framework[3]:

$$\texttt{JOIN_DLL}(hd,p,tl,n) \equiv \exists x,y \;.\; \texttt{DLL}(hd,p,x,y) * \texttt{DLL}(x,y,tl,n)$$

An extension to handle such definitions is possible, by introducing a special "concatenation" operation at the level of tree automata, which is considered for further investigation. □

Definition 4. *Given a system $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ of inductive definitions, a rule $R_{i,j}(x_{i,1},\ldots,x_{i,k}) \equiv \exists \mathbf{z} \;.\; \Sigma * P_{i_1}(\mathbf{y}_1) * \ldots * P_{i_m}(\mathbf{y}_m) \wedge \Pi$ is* connected *if and only if there exists a formal parameter $x_{i,\ell}$, $1 \leq \ell \leq k$ such that:*

- *$x_{i,\ell}$ is a root of Σ, and*
- *for each $j = 1,\ldots,m$, there exists $0 \leq s < |\mathbf{y}_j|$ such that $(\mathbf{y}_j)_s$ is $\Sigma \wedge \Pi$-reachable from $x_{i,\ell}$, and $x_{i_j,s}$ is a root of the head of each rule of P_{i_j}.*

An inductive system is said to be *connected* if all its rules are connected. In the rest of this section, we consider only connected systems. This condition is necessary, since later on it is proved that entailment between predicates of disconnected systems is undecidable, in general (Thm. 2 in Sec. 5). Notice that the DLL and TLL systems from Examples 1 and 2 are both connected.

2.5 Tree Automata

A (finite, non-deterministic, bottom-up) *tree automaton* (abbreviated as TA in the following) is a quadruple $A = \langle Q, \Sigma, \Delta, F \rangle$, where Σ is a finite alphabet, Q is a finite set of *states*, $F \subseteq Q$ is a set of *final states*, Σ is an alphabet, and Δ is a set of transition rules of the form $\sigma(q_1,\ldots,q_n) \to q$, for $\sigma \in \Sigma$, and $q, q_1, \ldots, q_n \in Q$. Given a tree automaton $A = \langle Q, \Sigma, \Delta, F \rangle$, for each rule $\rho = (\sigma(q_1,\ldots,q_n) \to q)$, we define its size as $|\rho| = n+1$. The size of the tree automaton is $|A| = \sum_{\rho \in \Delta} |\rho|$. A *run* of A over a tree $t : \mathbb{N}^* \rightharpoonup_{fin} \Sigma$ is a function $\pi : dom(t) \to Q$ such that, for each node $p \in dom(t)$, where $q = \pi(p)$, if $q_i = \pi(p.i)$ for $1 \leq i \leq n$, then Δ has a rule $(t(p))(q_1,\ldots,q_n) \to q$. We write $t \stackrel{\pi}{\Longrightarrow} q$ to denote that π is a run of A over t such that $\pi(\varepsilon) = q$. We use $t \Longrightarrow q$ to denote that $t \stackrel{\pi}{\Longrightarrow} q$ for some run π. The *language* of a A is defined as $\mathcal{L}(A) = \{t \mid \exists q \in F,\; t \Longrightarrow q\}$.

3 Tiles and Semantics of Inductive Definitions

A *tile* is a tuple $T = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1} \rangle$, for some $d \geq 0$, where φ is a basic SL formula and each $\mathbf{x}_i$ is a tuple of pairwise distinct variables, called a *port*, such that $\mathbf{x}_i \cap \mathbf{x}_j = \emptyset$, for all $-1 \leq i < j < d$, and $\mathbf{x}_{-1} \cup \mathbf{x}_0 \cup \ldots \cup \mathbf{x}_{d-1} \subseteq FV(\varphi)$. The set of all tiles is denoted

[3] However, defining concatenation of trees with linked leaves is possible when the root is allocated: $\texttt{JOIN_TLL}(root, ll, lr) \equiv \exists x,y,z \;.\; root \mapsto (x,y,\mathbf{nil}) * \texttt{TLL}(x,ll,z) * \texttt{TLL}(y,z,lr)$.

by $\mathcal{T}$. The variables from $\mathbf{x}_{-1}$ are said to be *incoming*, the ones from $\mathbf{x}_0, \ldots, \mathbf{x}_{d-1}$, are said to be *outgoing*, and the ones from $\mathbf{par}(T) = FV(\varphi) \setminus (\mathbf{x}_{-1} \cup \mathbf{x}_0 \cup \ldots \cup \mathbf{x}_{d-1})$ are called *parameters*. The *arity* of a tile $T = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1} \rangle$ is the number of outgoing ports, denoted by $\#(T) = d$. In the following, we denote $\mathbf{form}(T) \equiv \varphi$ and $\mathbf{port}_i(T) \equiv \mathbf{x}_i$, for all $i = -1, 0, \ldots, d-1$.

Given two tiles $T_1 = \langle \varphi_1, \mathbf{x}^1_{-1}, \mathbf{x}^1_0, \ldots, \mathbf{x}^1_{d-1} \rangle$ and $T_2 = \langle \varphi_2, \mathbf{x}^2_{-1}, \mathbf{x}^2_0, \ldots, \mathbf{x}^2_{e-1} \rangle$ such that $FV(\varphi_1) \cap FV(\varphi_2) = \emptyset$, for some $0 \leq i \leq d-1$ such that $|\mathbf{x}^1_i| = |\mathbf{x}^2_{-1}|$, we define their *i-composition*:

$$T_1 \circledast_i T_2 = \langle \psi, \mathbf{x}^1_{-1}, \mathbf{x}^1_0, \ldots \mathbf{x}^1_{i-1}, \mathbf{x}^2_0, \ldots, \mathbf{x}^2_{e-1}, \mathbf{x}^1_{i+1}, \ldots, \mathbf{x}^1_{d-1} \rangle$$
$$\text{where } \psi \equiv \exists \mathbf{x}^1_i \exists \mathbf{x}^2_{-1} \,.\, \varphi_1 * \varphi_2 \wedge \mathbf{x}^1_i = \mathbf{x}^2_{-1}.$$

For a tree position $q \in \mathbb{N}^*$ and a tile $T = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1} \rangle$, we denote by T^q the tile obtained by renaming every free variable $x \in \mathbf{x}_{-1} \cup \mathbf{x}_0 \cup \ldots \cup \mathbf{x}_{d-1}$ by x^q. Note that the parameters $x \in \mathbf{par}(T)$ are not changed by this renaming.

Definition 5. *A* tiled tree *is a tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}$ such that, for all positions $p \in dom(t)$, the following holds:*

- *$\#_t(p) = \#(t(p))$, i.e. the arity of p equals the arity of its label in t, and*
- *for all $0 \leq i < \#_t(p)$, the d-composition $t(p)^p \circledast_i t(p.i)^{p.i}$ is defined.*

A tiled tree t corresponds to a tile defined inductively, for any $p \in dom(t)$, as:

$$\Phi(t, p) = t(p)^p \circledast_0 \Phi(t, p.0) \circledast_1 \Phi(t, p.1) \ \ldots \ \circledast_{\#(p)-1} \Phi(t, p.(\#_t(p) - 1)).$$

The tile $\Phi(t, \varepsilon)$ is said to be the *characteristic tile* of t and it is denoted in the following as $\Phi(t)$. It can be easily shown, by induction on the structure of the tiled tree t, that $\Phi(t) = \langle \psi, \mathbf{x}_{-1} \rangle$, for an SL formula ψ and an incoming port $\mathbf{x}_{-1}$. In other words, $\Phi(t)$ has no outgoing ports. In this case, we write $S \models \Phi(t)$ for $S \models \psi$.

Given an inductive system $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ of the form (1), for each rule $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{u} \,.\, \Sigma * P_{i_0}(\mathbf{y}_0) * \ldots * P_{i_{n-1}}(\mathbf{y}_{d-1}) \ \wedge \ \Pi$, where $\mathbf{x} \cap \mathbf{u} = \emptyset$, we define a tile $T_{i,j} = \langle \varphi, \mathbf{x}, \mathbf{z}_0, \ldots, \mathbf{z}_{d-1} \rangle$, where $\mathbf{z}_i$ are disjoint from $\mathbf{y}_j$, for all $0 \leq i, j < d$, and $\varphi = \exists \mathbf{u} \,.\, \Sigma \wedge \Pi \ \wedge \ \bigwedge_{i=0}^{d-1} \mathbf{z}_i = \mathbf{y}_i$.

Definition 6. *Let $\mathcal{P} = \left\{P_i \equiv |_{j=1}^{m_i} R_{i,j}\right\}_{i=1}^{n}$ be an inductive system (1). An* unfolding tree *of $\mathcal{P}$ is a tiled tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}$ such that:*

- *$t(\varepsilon) = T_{i,j}$, for some $1 \leq j \leq m_i$, and*
- *for every position $p \in dom(t)$, if $t(p) = T_{i,j}$ and $\mathbf{tail}(R_{i,j}) = \langle P_{i_0}, \ldots, P_{i_{d-1}} \rangle$, where $d = \#(T_{i,j})$, then for all $0 \leq j < d$, we have $t(p.j) = T_{i_j,k}$, for some $1 \leq k \leq m_{i_j}$.*

If $t(\varepsilon) = T_{i,j}$, for some $1 \leq j \leq m_i$, we say that t is i-rooted.

We denote by $\mathcal{T}_i(\mathcal{P})$ the set of i-rooted unfolding trees of $\mathcal{P}$. The semantics of a predicate $P_i(\mathbf{x}) \in \mathcal{P}$ is defined as follows: $S \models P_i(\mathbf{x}) \iff S \models \Phi(t)$, for some $t \in \mathcal{T}_i(\mathcal{P})$. Given an inductive system $\mathcal{P}$ and two predicates $P_i(x_1, \ldots, x_n)$ and $P_j(y_1, \ldots, y_n)$ of $\mathcal{P}$, with the same number of formal parameters n, and a tuple of variables $\mathbf{x}$, where $|\mathbf{x}| = n$, the *entailment problem* is defined as follows: $P_i(\mathbf{x}) \models_{\mathcal{P}} P_j(\mathbf{x}) \ : \ \forall S \,.\, S \models P_i(\mathbf{x}) \Rightarrow S \models P_j(\mathbf{x})$. A *rooted* inductive system $\langle \mathcal{P}, P_i \rangle$ is an inductive system $\mathcal{P}$ with a designated predicate $P_i \in \mathcal{P}$. Two rooted systems $\langle \mathcal{P}, P_i \rangle$ and $\langle Q, Q_j \rangle$ are said to be *equivalent* if and only if $P_i \models_{\mathcal{P} \cup Q} Q_j$ and $Q_j \models_{\mathcal{P} \cup Q} P_i$.

Definition 7. *A rooted inductive system $\langle\{P_1,\ldots,P_n\},P_i\rangle$ is said to be* local *if and only if for each unfolding tree $t \in \mathcal{T}_i(\mathcal{P})$ and each state $S \models \Phi(t)$, S is local.*

The *locality problem* asks, given an inductive system $\mathcal{P} = \{P_1,\ldots,P_n\}$, and an index $i = 1,\ldots,n$, whether the rooted system $\langle\mathcal{P},P_i\rangle$ is local. Our method for deciding entailment problems of the form $P_i \models_{\mathcal{P}} P_j$ is shown to be sound and complete provided that both $\langle\mathcal{P},P_i\rangle$ and $\langle\mathcal{P},P_j\rangle$ are local. Otherwise, if one system is not local, our method is sound, i.e. the algorithm returns "yes" only if the entailment holds. We describe further a canonization procedure (Sec. 4.1) which performs a sufficient locality test on the system, prior to the encoding of the inductive system as a tree automaton.

3.1 Canonical Tiles

This section defines a class of *canonically tiled trees* (or, to be short, *canonical trees*) intended to reduce the number of ways in which a given local state can be encoded. Moreover, it is shown that despite a local state can still be described by several different canonical trees, such trees must be in a rotation relation called a *canonical rotation*. This fact is the basis of the completeness argument of our method. By defining a closure of tree automata (accepting canonical trees of states described by a given inductive predicate) under canonical rotations, we ensure that any entailment between two local inductive predicates can be reduced to a language inclusion problem. The case of non-local states is subsequently dealt with in a sound (yet, in general, not complete) way in the next subsection (Sec. 3.2).

A tile $T = \langle\varphi,\mathbf{x}_{-1},\mathbf{x}_0,\ldots,\mathbf{x}_{d-1}\rangle$ is said to be a *singleton* if φ is of one of the forms:

1. $\exists z\,.\, z \mapsto (y_0,\ldots,y_{m-1}) \wedge \Pi$ or
2. $z \mapsto (y_0,\ldots,y_{m-1}) \wedge \Pi$ and $z \in \mathbf{par}(T)$

and the following holds:

- For all $-1 \leq i < d$, $\mathbf{x}_i \cap (\{z,\mathbf{nil}\} \cup \mathbf{par}(T)) = \emptyset$, i.e. neither the incoming nor the outgoing tuples of variables contain z, **nil**, or parameters.
- For each $0 \leq j < m$, exactly one of the following holds: either (i) $y_j \equiv z$, (ii) $y_j \equiv \mathbf{nil}$, (iii) $y_j \in \mathbf{par}(T)$, or (iv) there exists a unique tuple of variables $\mathbf{x}_i$, $-1 \leq i < d$, such that y_j occurs in $\mathbf{x}_i$.
- For each outgoing tuple of variables $\mathbf{x}_i$, $0 \leq i < d$, there exists a variable $y_j \notin \{z,\mathbf{nil}\} \cup \mathbf{par}(T)$, for some $0 \leq j < m$, such that y_j occurs in $\mathbf{x}_i$.

We denote the spatial formula of a singleton tile as $(\exists z)\ z \mapsto (y_0,\ldots,y_{m-1})$ in order to account for the optional quantification of the allocated variable z. A singleton tile $T = \langle(\exists z)\ z \mapsto (y_0,\ldots,y_{m-1}) \wedge \Pi,\ \mathbf{x}_{-1},\mathbf{x}_0,\ldots,\mathbf{x}_{d-1}\rangle$ is said to be *canonical* if, moreover, for all $-1 \leq i < d$, $\mathbf{x}_i$ can be factorized as $\mathbf{x}_i \equiv \mathbf{x}_i^{fw} \cdot \mathbf{x}_i^{bw}$ such that:

1. $\mathbf{x}_{-1}^{bw} = \langle y_{h_0},\ldots,y_{h_k}\rangle$, for some ordered sequence $0 \leq h_0 < \ldots < h_k < m$, i.e. the incoming tuple is ordered by the selectors referencing its elements.
2. For all $0 \leq i < d$, $\mathbf{x}_i^{fw} \equiv \langle y_{j_0},\ldots,y_{j_{k_i}}\rangle$, for some ordered sequence $0 \leq j_0 < \ldots < j_k < m$, i.e. each outgoing tuple is ordered by the selectors referencing its elements.
3. For all $0 \leq i,j < d$, if $(\mathbf{x}_i^{fw})_0 \equiv y_p$ and $(\mathbf{x}_j^{fw})_0 \equiv y_q$, for some $0 \leq p < q < m$, then $i < j$, i.e. outgoing tuples are ordered by the selector referencing the first element.
4. $(\mathbf{x}_{-1}^{fw} \cup \mathbf{x}_0^{bw} \cup \ldots \cup \mathbf{x}_{d-1}^{bw}) \cap \{y_0,\ldots,y_{m-1}\} = \emptyset$ and $\Pi \equiv \mathbf{x}_{-1}^{fw} = z \wedge \bigwedge_{i=0}^{d-1} \mathbf{x}_i^{bw} = z$.

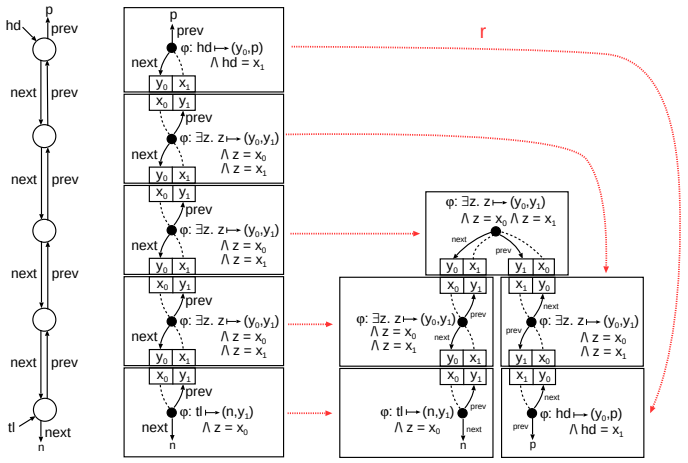

Fig. 4. An example of a DLL with two of its canonical trees (related by a canonical rotation r).

Given a canonical tile $T = \langle (\exists z)\ z \mapsto (y_0,\ldots,y_{m-1}) \wedge \Pi,\ \mathbf{x}_{-1},\mathbf{x}_0,\ldots,\mathbf{x}_{d-1}\rangle$, the decomposition of each tuple $\mathbf{x}_i$ into $\mathbf{x}_i^{fw}$ and $\mathbf{x}_i^{bw}$ is unique. This is because $\mathbf{x}_i^{fw}$ contains all referenced variables y_j that occur in an output port $\mathbf{x}_i$ (since $\mathbf{x}_i^{bw} \cap \{y_0,\ldots,y_{m-1}\} = \emptyset$) in the same order in which they are referenced in $z \mapsto (y_0,\ldots,y_{m-1})$. The same holds for the input port $\mathbf{x}_{-1}$, with the roles of $\mathbf{x}_{-1}^{fw}$ and $\mathbf{x}_{-1}^{bw}$ swapped. As each outgoing port must contain at least one such variable, by the definition of singleton tiles, we also have that $\mathbf{x}_i^{fw} \neq \emptyset$, for all $0 \leq i < d$. We denote further by $\mathbf{port}_i^{fw}(T)$ and $\mathbf{port}_i^{bw}(T)$ the tuples $\mathbf{x}_i^{fw}$ and $\mathbf{x}_i^{bw}$, resp., for all $i = -1,\ldots,d-1$. The set of canonical tiles is denoted as $\mathcal{T}^c$.

Definition 8. *A tiled tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$ is said to be* canonical *if and only if, for any $p \in dom(t)$ and each $0 \leq i < \#_t(p)$, we have $|\mathbf{port}_i^{fw}(t(p))| = |\mathbf{port}_{-1}^{fw}(t(p.i))|$ and $|\mathbf{port}_i^{bw}(t(p))| = |\mathbf{port}_{-1}^{bw}(t(p.i))|$.*

Example 3 (cont. of Ex. 1). To illustrate the notion of canonical trees, Fig. 4 shows two canonical trees for a given DLL. The tiles are depicted as big rectangles containing the appropriate basic formula as well the input and output ports. In all ports, the first variable is in the forward and the second in the backward part.

We will now work towards the notion of canonical rotations. For that, we first give one technical definition. Given a canonical tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$, and a state $S = \langle s,h\rangle$, let $u : dom(t) \to dom(h)$ be an arbitrary tree labeled with allocated locations from S. For each position $p \in dom(t)$ its *explicit neighbourhood* with respect to t and u is the state $S_{\langle t,u,p\rangle} = \langle s_p,h_p\rangle$ defined as follows:

- $h_p = \overline{h}$, and

– $s_p(x) = \begin{cases} u(p) & \text{if } x \in \mathbf{port}_{-1}^{fw}(t(p)) \cup \mathbf{port}_0^{bw}(t(p)) \cup \ldots \cup \mathbf{port}_{\#_t(p)-1}^{bw}(t(p)) \\ u(p.i) & \text{if } x \in \mathbf{port}_i^{fw}(t(p)), \text{ for some } 0 \leq i < \#_t(p) \\ u(p.(-1)) & \text{if } p \neq \varepsilon \text{ and } x \in \mathbf{port}_{-1}^{bw}(t(p)) \\ \overline{s}(x) & \text{otherwise} \end{cases}$

where $S_{\langle u(p) \rangle} = \langle \overline{s}, \overline{h} \rangle$ is the neighbourhood of the location $u(p)$ in S.

An important property of canonical trees is that each state that is a model of the characteristic tile $\Phi(t)$ of a canonical tree t (i.e. $S \models \Phi(t)$) can be uniquely described by a local spanning tree $u : dom(t) \to Loc$, which has the same structure as t, i.e. $dom(u) = dom(t)$. An immediate consequence is that any two models of $\Phi(t)$ differ only by a renaming of the allocated locations, i.e. they are identical up to isomorphism.

Lemma 3. *Let $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$ be a canonical tree, and let $S = \langle s, h \rangle$ be a state. Then $S \models \Phi(t)$ iff there exists a local spanning tree $u : dom(t) \to dom(h)$ such that, for all $p \in dom(t)$:*

1. $|\mathbf{port}_i^{fw}(t(p))| = \|\{k \in Sel \mid u(p) \xrightarrow{k}_S u(p.i)\}\|$, *for all* $0 \leq i < \#_t(p)$.
2. $|\mathbf{port}_{-1}^{bw}(t(p))| = \|\{k \in Sel \mid u(p) \xrightarrow{k}_S u(p.(-1))\}\|$ *if* $p \neq \varepsilon$.
3. $S_{\langle t,u,p \rangle} \models \mathbf{form}(t(p))$.

Proof. "$\Rightarrow$" By induction on the structure of t. For the base case $\#_t(\varepsilon) = 0$, i.e. $dom(t) = \{\varepsilon\}$, we have $t(\varepsilon) = \langle \mathbf{form}(t(\varepsilon), \mathbf{port}_{-1}(t(\varepsilon)) \rangle$ and therefore $S \models \Phi(t)$ if and only if $S \models \mathbf{form}(t(\varepsilon))$. Since $t(\varepsilon)$ is a canonical tile, only one location is allocated in $dom(h)$, i.e. $dom(h) = \{\ell\}$, for some $\ell \in Loc$. We define $u = \{(\varepsilon, \ell)\}$. It is immediate to check that u is a local spanning tree of S, and that $S_{\langle t,u,\varepsilon \rangle} = S \models \mathbf{form}(t(\varepsilon))$ – points (1) and (2) are vacuously true.

For the induction step $\#_t(\varepsilon) = d > 0$, we have:

$$\Phi(t, \varepsilon) = t(\varepsilon)^{\varepsilon} \circledast_0 \Phi(t, 0) \ldots \circledast_{d-1} \Phi(t, d-1)$$

Since $S \models \Phi(t, \varepsilon)$, there exist states $S_i = \langle s_i, h_i \rangle$, for each $-1 \leq i < d$, such that $S_{-1} \models t(\varepsilon)^{\varepsilon}$ and $S_i \models \Phi(t, i)$, for all $0 \leq i < d$, and moreover, $S = S_{-1} \uplus S_0 \uplus \ldots \uplus S_{d-1}$. By the induction hypothesis, for each $0 \leq i < d$ there exists a local spanning tree $u_i : dom(t_{|_i}) \to dom(h_i)$, meeting conditions (1), (2) and (3). Since $S_{-1} \models t(\varepsilon)^{\varepsilon}$, it must be that $dom(h_{-1}) = \{\ell\}$ for some location $\ell \in Loc$. We define u as follows:

$$\begin{array}{l} u(\varepsilon) = \ell \\ u(i.q) = u_i(q) \text{ for all } q \in dom(u_i) \end{array}$$

To prove that u is a spanning tree of S, let $p \in dom(u)$ be a position, and $0 \leq i < \#_u(p)$ be a direction. We distinguish two cases:

– if $p = \varepsilon$, then $\mathbf{port}_i^{fw}(t(p)) \neq \emptyset$, hence there exists an edge $\ell \xrightarrow{k}_{S_{-1}} u_i(\varepsilon)$. Then $u(p) \xrightarrow{k}_S u(p.i)$ as well.
– if $p = j.q$, and $q \in dom(u_j)$, for some $0 \leq j < d$, by the induction hypothesis, there exists $k \in Sel$, such that $u(p) = u_j(q) \xrightarrow{k}_{S_j} u_j(q.i) = u(p.i)$, i.e. $u(p) \xrightarrow{k}_S u(p.i)$.

To prove that u is a local spanning tree of S, let $\kappa \xrightarrow{k}_S \kappa'$ be an edge, for some $\kappa, \kappa' \in dom(h)$. Since $dom(h)$ is a disjoint union of $dom(h_{-1}) = \{\ell\}, dom(h_0), \ldots, dom(h_{d-1})$, we distinguish several cases:

- if $\kappa \in dom(h_{-1})$ and $\kappa' \in dom(h_i)$, for some $0 \leq i < d$, then $u(p) = \kappa = \ell$ and $u(i) = u_i(\varepsilon) = \kappa'$ is the only possibility – due to the strict semantics of SL, it is not possible to define an edge between ℓ and an location $\kappa' = u_i(q)$, unless $q = \varepsilon$.
- if $\kappa, \kappa' \in dom(h_i)$, for some $0 \leq i < d$, then, by the induction hypothesis, there exists $q \in dom(u_i)$ and $0 \leq j < \#_{u_i}(q)$ such that $\kappa = u_i(q) = u(i.q)$ and $\kappa' = u_i(q.j) = u(i.q.j)$.
- if $\kappa \in dom(h_i)$ and $\kappa' \in dom(h_j)$, for some $0 \leq i, j < d$, then we reach a contradiction with the strict semantics of SL – since there is no equality between output ports in $t(\varepsilon)$, it is not possible to define an edge between a location from $dom(h_i)$ and $dom(h_j)$,

The proofs of points (1), (2) and (3) are by the strict semantics of SL.

"$\Leftarrow$" By induction on the structure of t. In the base case $\#_t(\varepsilon) = 0$, i.e. $dom(t) = \{\varepsilon\}$, then $t(\varepsilon) = \langle \mathbf{form}(t(\varepsilon), \mathbf{port}_{-1}(t(\varepsilon)) \rangle$. We have $S = S_{\langle t,u,\varepsilon \rangle} \models \mathbf{form}(t(\varepsilon))$, by point (3), hence $S \models \Phi(t)$.

For the induction step $\#_t(\varepsilon) = d > 0$, let $u_i = u_{|_i}$, for all $0 \leq i < d$. Since u is a bijective function, we have $img(u_i) \cap img(u_j) = \emptyset$, for all $0 \leq i < j < d$. For all $-1 \leq i < d$, we define $S_i = \langle s_i, h_i \rangle$, where:

$$h_i(\ell) = \begin{cases} h(\ell) \text{ if } \ell \in img(u_i) \\ \bot \quad \text{otherwise} \end{cases} \qquad \text{for all } \ell \in Loc, \text{ if } i \geq 0$$

$$h_{-1}(\ell) = \begin{cases} h(\ell) \text{ if } u(\varepsilon) = \ell \\ \bot \quad \text{otherwise} \end{cases} \qquad \text{for all } \ell \in Loc$$

$$s_i(x) = \begin{cases} s(x) \text{ if } s(x) \in Img(h_i) \\ \bot \quad \text{otherwise} \end{cases} \qquad \text{for all } x \in Var$$

It is not hard to show that $S = S_{-1} \uplus S_0 \uplus \ldots S_{d-1}$. Since u is a local spanning tree of S, u_i is a local spanning tree for S_i, for all $0 \leq i < d$, and conditions (1), (2) and (3) hold for S_i and u_i, respectively. By the induction hypothesis, we have $S_i \models \Phi(t_{|_i}, \varepsilon)$. By points (1), (2) and (3) moreover, we have that $S_{-1} \models \mathbf{form}(t(\varepsilon))$. Hence $S \models \Phi(t)$. $\square$

The following definition is a refinement of Def. 2. The change in the structure of the tree is mirrored by a change in the structure of the canonical tiles labeling the tree.

Definition 9. *Given two canonical trees $t, u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$, we say that u is a* canonical rotation *of t, denoted $t \sim^c_r u$, if and only if $r : dom(t) \to dom(u)$ is a bijective function, and for all $p \in dom(t)$, there exists a substitution $\sigma_p : Var \rightharpoonup_{fin} Var$ such that $\mathbf{form}(t(p))[\sigma_p] \equiv \mathbf{form}(u(r(p)))$, and for all $0 \leq i < \#_t(p)$, there exists $j \in \mathcal{D}(u)$ such that $r(p.i) = r(p).j$, and:*

$$\mathbf{port}_i^{fw}(t(p))[\sigma_p] \equiv \textbf{if } j \geq 0 \textbf{ then } \mathbf{port}_j^{fw}(u(r(p))) \textbf{ else } \mathbf{port}_{-1}^{bw}(u(r(p)))$$
$$\mathbf{port}_i^{bw}(t(p))[\sigma_p] \equiv \textbf{if } j \geq 0 \textbf{ then } \mathbf{port}_j^{bw}(u(r(p))) \textbf{ else } \mathbf{port}_{-1}^{fw}(u(r(p)))$$

Example 4 (cont. of Ex. 3). The notion of canonical rotation is illustrated by the canonical rotation r relating the two canonical trees of a DLL shown in Fig. 4. In its case, the variable substitutions are simply the identity in each node. Note, in particular, that when the tile 0 of the left tree (i.e., the second one from the top) gets rotated to the tile 1 of the right tree (i.e., the right successor of the root), the input and output ports get swapped and so do their forward and backward parts.

The following lemma is the key for proving completeness of our entailment checking for local inductive systems: if a (local) state is a model of the characteristic tiles of two different canonical tree, then these trees must be related by canonical rotation.

Lemma 4. *Let $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$ be a canonical tree and $S = \langle s, h \rangle$ be a state such that $S \models \Phi(t)$. Then, for any canonical tree $u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^c$, we have $S \models \Phi(u)$ iff $t \sim^c u$.*

Proof. If $S \models \Phi(t)$, there exists a local spanning tree $t_s : dom(t) \to dom(h)$ of S, that meets the three conditions of Lemma 3.

"$\Rightarrow$" If $S \models \Phi(u)$ there exists a local spanning tree $u_s : dom(u) \to dom(h)$ of S, that meets the three conditions of Lemma 3. Since t_s and u_s are spanning trees of S, by Lemma 2, we have $t_s \sim_{u_s^{-1} \circ t_s} u_s$. Let $r = u_s^{-1} \circ t_s$ from now on. Clearly, for any $p \in dom(t)$, we have $t_s(p) = u_s(r(p))$. Since $dom(t) = dom(t_s)$ and $dom(u) = dom(u_s)$, we have $t \sim_r u$ as well.

We further need to prove the three points of Def. 9 in order to show that $t \sim_r^c u$. Let $p \in dom(t)$ be an arbitrary location. For the first two points, let $0 \leq i < \#_t(p)$ be a direction. We have $t_s(p.i) = u_s(r(p.i)) = u_s(r(p).j)$, for some $j \in \mathcal{D}(u)$. We distinguish two cases:

– if $j \geq 0$, we compute:

$$\begin{aligned}
|\mathbf{port}_i^{fw}(t(p))| &= \|\{k \in Sel \mid t_s(p) \xrightarrow{k}_S t_s(p.i)\}\| \\
&= \|\{k \in Sel \mid u_s(r(p)) \xrightarrow{k}_S u_s(r(p).j)\}\| \\
&= |\mathbf{port}_j^{fw}(u(r(p)))|
\end{aligned}$$

$$\begin{aligned}
|\mathbf{port}_i^{bw}(t(p))| &= |\mathbf{port}_{-1}^{bw}(t(p.i))| \text{ (by Def. 8)} \\
&= \|\{k \in Sel \mid t_s(p.i) \xrightarrow{k}_S t_s(p)\}\| \\
&= \|\{k \in Sel \mid u_s(r(p).j) \xrightarrow{k}_S u_s(r(p))\}\| \\
&= |\mathbf{port}_j^{bw}(u(r(p)))|
\end{aligned}$$

– if $j = -1$, we compute:

$$\begin{aligned}|\mathbf{port}_i^{fw}(t(p))| &= \|\{k \in Sel \mid u_s(r(p)) \xrightarrow{k}_S u_s(r(p).(-1))\}\| \\ &= |\mathbf{port}_{-1}^{bw}(u(r(p)))|\end{aligned}$$

$$\begin{aligned}|\mathbf{port}_i^{bw}(t(p))| &= \|\{k \in Sel \mid u_s(r(p).(-1)) \xrightarrow{k}_S u_s(r(p))\}\| \\ &= \|\{k \in Sel \mid u_s(r(p).(-1)) \xrightarrow{k}_S u_s(r(p).(-1).j')\}\| \\ &\quad \text{for some } j' \geq 0 \text{ such that } r(p) = r(p).(-1).j' \\ &= |\mathbf{port}_{j'}^{fw}(u(r(p).(-1)))| \\ &= |\mathbf{port}_{-1}^{fw}(u(r(p)))| \text{ (by Def. 8)}\end{aligned}$$

Since all variables are pairwise distinct in $\mathbf{port}_i(t(p))$ and $\mathbf{port}_j(u(r(p)))$, respectively, one can define a substitution σ_p meeting the conditions of the first two points of Def. 9.

For the third point of Def. 9, by Lemma 3, we have that $S_{\langle t,t_s,p\rangle} \models \mathbf{form}(t(p))$ and $S_{\langle u,u_s,r(p)\rangle} \models \mathbf{form}(u(r(p)))$, where $S_{\langle t,t_s,p\rangle} = \langle s_t, h_t\rangle$ is the explicit neighbourhood of p w.r.t t and t_s, and $S_{\langle u,u_s,p\rangle} = \langle s_u, h_u\rangle$ is the explicit neighbourhood of $r(p)$ w.r.t u and u_s. Since $t(p)$ and $u(r(p))$ are canonical tiles, we have:

$$\begin{aligned}\mathbf{form}(t(p)) &\equiv (\exists z)\ z \mapsto (y_0, \ldots, y_{m-1}) \wedge \Pi_t \\ \mathbf{form}(u(r(p))) &\equiv (\exists w)\ w \mapsto (v_0, \ldots, v_{n-1}) \wedge \Pi_u\end{aligned}$$

where

$$\begin{aligned}\Pi_t &\equiv \mathbf{port}_{-1}^{fw}(t(p)) = z \wedge \textstyle\bigwedge_{i=0}^{\#_t(p)-1} \mathbf{port}_i^{bw}(t(p)) = z \\ \Pi_u &\equiv \mathbf{port}_{-1}^{fw}(u(r(p))) = w \wedge \textstyle\bigwedge_{i=0}^{\#_u(r(p))-1} \mathbf{port}_i^{bw}(u(r(p))) = w\end{aligned}$$

We extend the substitution σ_p to $\sigma_p[z \leftarrow w]$ in the following. It is not hard to check that $s_t \circ \sigma_p = s_u$ and $h_t = h_u$. Hence $S_{\langle u,u_s,p\rangle} \models \mathbf{form}(t(p))[\sigma_p]$. Since both $\mathbf{form}(t(p))[\sigma_p]$ and $\mathbf{form}(u(r(p)))$ have the same model, by the definition of the (strict) semantics of SL, it follows that the numbers of edges are the same, i.e. $m = n$, and either both z and w are quantified, or they are free. Thus, we obtain that $\mathbf{form}(t(p))[\sigma_p] \equiv \mathbf{form}(u(r(p)))$.

"$\Leftarrow$" It $t \sim^c u$, there exists a bijective function $r : dom(t) \to dom(u)$, meeting the conditions of Def. 9. Let $u_s = t_s \circ r^{-1}$ be a tree. Since t_s and r are bijective, then also u_s is bijective, and $dom(u_s) = dom(u)$. To prove that $S \models \Phi(u)$, it is enough to show that u_s is a local spanning tree of S, meeting the three conditions of Lemma 3.

To show that u_s is a local spanning tree of S, let $p \in dom(u_s)$ and $i \in \mathcal{D}_+(u_s)$ such that $p.i \in dom(u_s)$. Since $t_s \sim_r u_s$, by Prop. 1, we have $u_s \sim_{r^{-1}} t_s$, hence there exists $j \in \mathcal{D}(t_s)$ such that $r^{-1}(p.i) = r^{-1}(p).j$. Since t_s is a spanning tree of S, there exists a selector $k \in Sel$ such that:

$$u_s(p) = t_s(r^{-1}(p)) \xrightarrow{k}_S t_s(r^{-1}(p).j) = t_s(r^{-1}(p.i)) = u_s(p.i)$$

Hence u_s is a spanning tree of S. In order to prove its locality, let $\ell \xrightarrow{k} \ell'$ be an edge of S, for some $\ell, \ell \in loc(S)$ and $k \in Sel$. Since t_s is a local spanning tree of S, there exist $p \in dom(t_s)$ and $i \in \mathcal{D}_+(t_s)$ such that $\ell = t_s(p) = u_s(r(p))$ and $\ell' = t_s(p.i) = u_s(r(p.i))$. Since $t_s \sim_r u_s$, there exists $j \in \mathcal{D}(u_s)$ such that $r(p.i) = r(p).j$. Hence we have $u_s(r(p)) = \ell \xrightarrow{k} \ell' = u_s(r(p).j)$, i.e. the edge is local w.r.t. u_s as well.

To prove point (1) of Lemma 3, let $p \in dom(u)$ be an arbitrary position, and $0 \leq i < \#_u(p)$ be a direction. Since $u_s \sim_{r^{-1}} t_s$, we have $r^{-1}(p.i) = r^{-1}(p).j$, for some $j \in \mathcal{D}(t)$. We compute:

$$\begin{aligned}\|\{k \in Sel \mid u_s(p) \xrightarrow{k}_S u_s(p.i)\}\| &= \|\{k \in Sel \mid t_s(r^{-1}(p)) \xrightarrow{k}_S t_s(r^{-1}(p.i)\}\| \\ &= \|\{k \in Sel \mid t_s(r^{-1}(p)) \xrightarrow{k}_S t_s(r^{-1}(p).j)\}\|\end{aligned}$$

We distinguish two cases:

- if $j \geq 0$, we have, by Lemma 3 applied to t and t_s:

$$\begin{aligned}\|\{k \in Sel \mid t_s(r^{-1}(p)) \xrightarrow{k}_S t_s(r^{-1}(p).j)\}\| &= |\mathbf{port}_j^{fw}(t(r^{-1}(p)))| \\ &= |\mathbf{port}_i^{fw}(u(p))|\end{aligned}$$

  The last equality is because $r(r^{-1}(p).j) = r(r^{-1}(p)).i$.
- if $j = -1$, there exists $d \in \mathcal{D}_+(t)$ such that $r^{-1}(p) = (r^{-1}(p).(-1)).d$. But then we have:

$$\begin{aligned}r((r^{-1}(p).(-1)).d) &= r(r^{-1}(p)) = p = (p.i).(-1) \\ &= r(r^{-1}(p).(-1)).(-1)\end{aligned}$$

  We compute, further:

$$\begin{aligned}\|\{k \in Sel \mid t_s(r^{-1}(p)) \xrightarrow{k}_S t_s(r^{-1}(p).j)\}\| &= |\mathbf{port}_{-1}^{bw}(t(r^{-1}(p)))| \\ &= |\mathbf{port}_d^{bw}(t(r^{-1}(p).(-1)))| \text{ (by Def. 8)} \\ &= |\mathbf{port}_{-1}^{fw}(u(p.i))| \text{ (by Def. 9)} \\ &= |\mathbf{port}_i^{fw}(u(p))| \text{ (by Def. 8)}\end{aligned}$$

For point (2) of Lemma 3, one applies a symmetrical argument. To prove point (3) of Lemma 3, let $S_{\langle t,t_s,r^{-1}(p)\rangle} = \langle s_t, h_t\rangle$ be the explicit neighbourhood of $r^{-1}(p)$ w.r.t t and t_s. It is not hard to check that $S_{\langle u,u_s,p\rangle} = \langle s_t \circ \sigma_p, h_t\rangle$, hence:

$$S_{\langle u,u_s,p\rangle} = \langle s_t \circ \sigma_p, h_t\rangle \models \mathbf{form}(t(r^{-1}(p)))[\sigma_p] \equiv \mathbf{form}(u(p))$$

The last equivalence is by Def. 9. □

3.2 Quasi-Canonical Tiles

This section introduces a generalization of canonical trees, denoted as *quasi-canonical trees*, that are used to define states with non-local edges. A singleton tile $T = \langle \varphi \wedge \Pi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1}\rangle$, for some $d \geq 0$, is said to be *quasi-canonical* if and only if $\mathbf{x}_i \equiv \mathbf{x}_i^{fw} \cdot \mathbf{x}_i^{bw} \cdot \mathbf{x}_i^{eq}$, for each $-1 \leq i < d$, Π is pure formula, and all of the following holds:

1. $\langle \varphi,\ \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{bw},\ \mathbf{x}_0^{fw} \cdot \mathbf{x}_0^{bw},\ \ldots,\ \mathbf{x}_{d-1}^{fw} \cdot \mathbf{x}_{d-1}^{bw}\rangle$ is a canonical tile.
2. For each $0 \leq i < |\mathbf{x}_{-1}^{eq}|$, either $(\mathbf{x}_{-1}^{eq})_i \in FV(\varphi)$ or $(\mathbf{x}_{-1}^{eq})_i =_\Pi (\mathbf{x}_k^{eq})_j$ for some unique indices $0 \leq k < d$ and $0 \leq j < |\mathbf{x}_k^{fw}|$.
3. For each $0 \leq k < d$ and each $0 \leq j < |\mathbf{x}_k^{eq}|$, either $(\mathbf{x}_k^{eq})_j \in FV(\varphi)$ or exactly one of the following holds: (i) $(\mathbf{x}_k^{eq})_j =_\Pi (\mathbf{x}_{-1}^{eq})_i$ for some unique index $0 \leq i < |\mathbf{x}_{-1}^{eq}|$. (ii) $(\mathbf{x}_k^{eq})_j =_\Pi (\mathbf{x}_r^{eq})_s$ for some unique indices $0 \leq r < d$ and $0 \leq s < |\mathbf{x}_r^{eq}|$.
4. For any two variables $x, y \in \mathbf{x}_{-1}^{eq} \cup \bigcup_{i=0}^{d-1} \mathbf{x}_i^{eq}$, $x =_\Pi y$ only in one of the cases above.

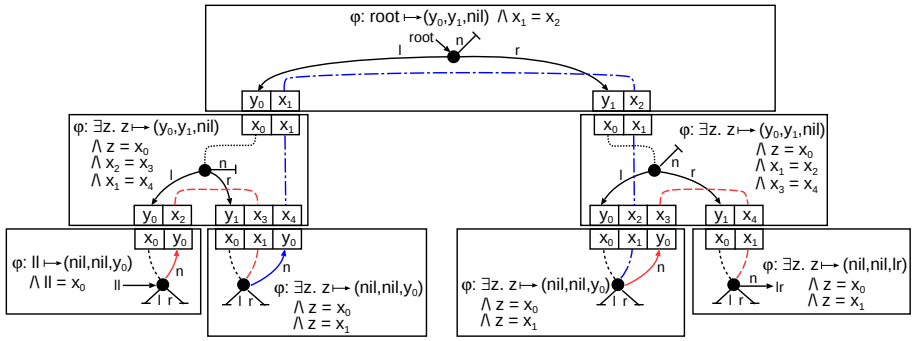

Fig. 5. An example of a quasi-canonically tiled tree for the tree with linked leaves from Fig. 3.

Notice that in a quasi-canonical tile $T = \langle (\exists z)\, z \mapsto (y_0, \ldots, y_{m-1}) \wedge \Pi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1} \rangle$, the free variables from the set $\{y_0, \ldots, y_{m-1}\} \setminus \{z, \mathbf{nil}\}$ are allowed to occur in some tuple $\mathbf{x}_i^{eq}$, for $-1 \leq i < d$, as well. These variables can be used to describe non-local edges of states because equality constraints between $\mathbf{x}_{-1}^{eq}$ and $\mathbf{x}_i^{eq}$, for $0 \leq i < d$, may span several tiles in a tree. In this case, the factorization of each port $\mathbf{x}_i$, $-1 \leq i < d$, into $\mathbf{x}_i^{fw}$, $\mathbf{x}_i^{bw}$, and $\mathbf{x}_i^{eq}$ is not uniquely determined by the distribution of the referenced variables $y_0, \ldots, y_{m-1}$. In practice (Sec.4.1) we use a conservative heuristic to define this factorization. We denote in the following $\mathbf{port}_i^{eq}(T) \equiv \mathbf{x}_i^{eq}$, for all $-1 \leq i < d$. The set of quasi-canonical tiles is denoted by $\mathcal{T}^{qc}$. The next definition of quasi-canonical trees extends Def. 8 to the case of quasi-canonical tiles.

Definition 10. *A tiled tree* $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$ *is* quasi-canonical *if and only if, for any* $p \in dom(t)$ *and each* $0 \leq i < \#_t(p)$, *such that* $t(p) = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{d-1} \rangle$, $t(p.i) = \langle \psi, \mathbf{y}_{-1}, \mathbf{y}_0, \ldots, \mathbf{y}_{e-1} \rangle$, $\mathbf{x}_i \equiv \mathbf{x}_i^{fw} \cdot \mathbf{x}_i^{bw} \cdot \mathbf{x}_i^{eq}$ *and* $\mathbf{y}_{-1} \equiv \mathbf{y}_{-1}^{fw} \cdot \mathbf{y}_{-1}^{bw} \cdot \mathbf{y}_{-1}^{eq}$, *we have* $|\mathbf{x}_i^{fw}| = |\mathbf{y}_{-1}^{fw}|$, $|\mathbf{x}_i^{bw}| = |\mathbf{y}_{-1}^{bw}|$, *and* $|\mathbf{x}_i^{eq}| = |\mathbf{y}_{-1}^{eq}|$.

Example 5 (cont. of Ex. 2). For an illustration of the notion of canonical trees, see Fig. 5, which shows a canonical tree for the TLL from Fig. 3. The figure uses the same notation as Fig. 4. In all the ports, the first variable is in the forward part, the backward part is empty, and the rest is the equality part.

For a quasi-canonical tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$, we define its *canonical projection* as $t^c : dom(t) \to \mathcal{T}^c$, such that, for all $p \in dom(t)$, we have:

$$t^c(p) = \langle \psi, \mathbf{port}_{-1}^{fw}(t(p)) \cdot \mathbf{port}_{-1}^{bw}(t(p)), \ldots, \mathbf{port}_{\#_t(p)-1}^{fw}(t(p)) \cdot \mathbf{port}_{\#_t(p)-1}^{bw}(t(p)) \rangle$$
$$\text{where } \psi \iff \exists \mathbf{port}_{-1}^{eq}(t(p)) \ldots \exists \mathbf{port}_{\#_t(p)-1}^{eq}(t(p)) \,.\, \mathbf{form}(t(p))$$

Intuitively, the canonical projection of a tree removes all equality variables from every tile in the tree. Since these variables are involved only in equality constraints, the above quantifiers $\exists \mathbf{port}_i^{eq}(t(p))$ can be eliminated. We can now define the notion of quasi-canonical rotation for quasi-canonical trees as an extension of Def. 9.

Definition 11. *Given two quasi-canonical trees* $t, u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$, *we say that* u *is a* quasi-canonical rotation *of* t, *denoted* $t \sim_r^{qc} u$, *if and only if* $r : dom(t) \to dom(u)$ *is*

a bijective function such that $t^c \sim_r^c u^c$ *and, for all* $p \in dom(t)$*, there exists a substitution* $\sigma_p : Var \rightharpoonup_{fin} Var$ *such that, for all* $0 \leq i < \#_t(p)$ *and all* $-1 \leq j < \#_t(p)$*, where* $r(p.i) = r(p).j$*, we have* $\mathbf{port}_i^{eq}(t(p))[\sigma] \equiv \mathbf{port}_j^{eq}(u(r(p)))$.

The increase in expressivity (i.e. the possibility of defining non-local edges) comes at the cost of a loss of completeness. The following lemma generalizes the necessity direction ($\Leftarrow$) of Lemma 4 for quasi-canonical tiles. Notice that the sufficiency ($\Rightarrow$) direction does not hold in general.

Lemma 5. *Let* $t, u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$ *be quasi-canonical trees such that* $t \sim^{qc} u$*. For all states* S*, if* $S \models \Phi(t)$*, then* $S \models \Phi(u)$.

Proof. Let S be a state such that $S \models \Phi(t)$, and $r : dom(t) \to dom(u)$ be a bijective function such that $t \sim_r^{qc} u$. Then $S \models \Phi(t^c)$, and since we have $t^c \sim_r^c u^c$ by Def. 11, then also $S \models \Phi(u^c)$, by Lemma 4. Also, for every $p \in dom(t)$, let $\sigma_p : Var \rightharpoonup_{fin} Var$ be the substitution from Def. 9.

Assume that $S \not\models \Phi(u)$. Then there exists a non-trivial path $p_0, \ldots, p_k \in dom(u)$, for some $k > 0$, and some variables $x_0 \in FV(\mathbf{form}(u(p_0))), \ldots, x_k \in FV(\mathbf{form}(u(p_k)))$, such that x_0 and x_k are allocated in $\mathbf{form}(u(p_0))$ and $\mathbf{form}(u(p_k))$, respectively, and for all $0 \leq i < k$, we have $x_i =_{\Pi_u} x_{i+1}$, where Π_u is the pure part of $\mathbf{form}(\Phi(u))$. But then $r^{-1}(p_0), \ldots, r^{-1}(p_k) \in dom(t)$ is a path in t, and the variables $\sigma_{p_0}^{-1}(x_0)$ and $\sigma_{p_k}^{-1}(x_k)$ are allocated in $t(r^{-1}(p_0))$ and $t(r^{-1}(p_k))$, respectively. Moreover, we have $\sigma_{p_i}^{-1}(x_i) =_{\Pi_t} \sigma_{p_{i+1}}^{-1}(x_{i+1})$, where Π_t is the pure part of $\mathbf{form}(\Phi(t))$. Hence $S \not\models \Phi(t)$, contradiction. □

4 From Inductive Definitions to Tree Automata

This section describes the reduction of an entailment problem, for given inductive systems, to a language inclusion problem between two tree automata. This reduction is the basis of our method for entailments checking. The complexity of the reduction algorithm is polynomial in the size of the input system.

4.1 Canonization

Given a rooted inductive system $\langle \mathcal{P}, P_i \rangle$, the canonization procedure described in this section will produce an equivalent rooted inductive system $\langle Q, Q_j \rangle$, such that each unfolding tree $t \in \mathcal{T}_j(Q)$ is canonical. For the latter inductive system, we define a TA recognizing its set of unfolding trees. For the rest of this section, let $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ be an inductive system of the form (1), where each rule is of the form $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \Sigma * P_{i_1}(\mathbf{y}_1) * \ldots * P_{i_m}(\mathbf{y}_m) \wedge \Pi$, $\mathbf{y}_1 \cup \ldots \cup \mathbf{y}_m \cup FV(\Pi) \subseteq \mathbf{x} \cup \mathbf{z}$. We recall that Σ is a non-empty spatial formula, and the pure formula Π can only specify equalities between (i) two existentially quantified variables from $\mathbf{z}$, (ii) an existentially quantified variable from $\mathbf{z}$ and a formal parameter from $\mathbf{x}$, or (iii) between an allocated parameter from $\mathbf{x}$ and some other formal parameter in $\mathbf{x}$. Notice that equality between two different allocated variables leads to unsatisfiable rules, therefore it will cause a pre-processing error.

Elimination of equalities. The first pre-processing step consists in eliminating equalities involving existentially quantified variables from the pure part Π of a rule. We use the equivalence relation $=_\Pi$ induced by Π on the variables in the rule. For each variable $x \in \mathbf{x} \cup \mathbf{z}$, its equivalence class $[x]_\Pi$ either (i) contains at most one formal parameter, or (ii) contains one or more formal parameters, out of which exactly one is allocated. We consider this parameter to be the representative of the equivalence class, and replace each occurrence of a variable in a rule by its representative. The result is an equivalent system of the same size as the original.

Let $\mathbf{alloc}(\Sigma)$ denote the set of allocated variables in Σ. For each equivalence class $[x]_\Pi$, where $x \in \mathbf{x} \cup \mathbf{z}$, we define its *representative* to be either one of the following:

- the unique formal parameter $x \in [x]_\Pi$, if $[x]_\Pi \cap \mathbf{x} \neq \emptyset$ and $[x]_\Pi \cap \mathbf{x} \cap \mathbf{alloc}(\Sigma) = \emptyset$,
- the unique allocated parameter $x \in [x]_\Pi \cap \mathbf{x} \cap \mathbf{alloc}(\Sigma)$, if $[x]_\Pi \cap \mathbf{x} \cap \mathbf{alloc}(\Sigma) \neq \emptyset$,
- the lexicographically minimal element $\mathit{minlex}([x]_\Pi)$ of $[x]_\Pi$, if $[x]_\Pi \cap \mathbf{x} = \emptyset$.

For a tuple of variables $\mathbf{y} = \langle y_1, \ldots, y_\ell \rangle$, let $[\mathbf{y}]_\Pi$ denote the tuple of representatives of the equivalence classes $[y_1]_\Pi, \ldots, [y_\ell]_\Pi$, respectively. Then we define the rule:

$$R^{=}_{i,j}(\mathbf{x}) \equiv \Sigma^{=} * P_{i_1}([\mathbf{y}_1]_\Pi) * \ldots * P_{i_m}([\mathbf{y}_m]_\Pi) \wedge \Pi^{=}$$

where $\Sigma^{=}$ is obtained from Σ by replacing each free variable $x \in FV(\Sigma)$ by the representative of $[x]_\Pi$, and $\Pi^{=} \equiv \bigwedge\{x = y \mid x,y \in \mathbf{x},\ x \in [y]_\Pi \cap \mathbf{alloc}(\Sigma)\}$ keeps only the equalities between formal parameters, one of which is allocated. It is not hard to check that any rooted system $\langle \mathcal{P}[R^{=}_{i,j}/R_{i,j}], P_k \rangle$, obtained by replacing $R_{i,j}$ with $R^{=}_{i,j}$, is equivalent to the original rooted system $\langle \mathcal{P}, P_k \rangle$, for all $k = 1, \ldots, n$. Moreover, the two systems have the same size.

Reduction to one points-to proposition per rule. The second step of the canonization procedure builds an inductive system in which the head of each rule consists of exactly one points-to atomic proposition of the form $x \mapsto (y_0, \ldots, y_\ell)$. This step succeeds under the assumption that $\mathcal{P}$ is a connected system, and no rule has an empty head (otherwise the system breaks the restrictions that we have introduced, and the computation is aborted). The result is an equivalent connected system Q, whose size is increased by at most a linear factor, i.e. $|Q| = O(|\mathcal{P}|)$. Algorithm 1 describes the splitting of rules into sets of rules with exactly one points-to proposition.

For each rule $R_{i,j}$, with $\mathbf{head}(R_{i,j}) \equiv \Sigma$, the procedure DEPTHFIRSTTRAVERSE(Σ, x) (line 8) performs a depth-first traversal of the spatial formula $\Sigma \equiv \bigstar_{i=1}^{s} x_i \mapsto (y_{i,1}, \ldots, y_{i,m_i})$ starting with a randomly chosen free variable $x \in FV(\Sigma)$, and builds an injective depth-first spanning tree $t : \mathbb{N}^* \rightharpoonup_{fin} AP(\Sigma)$ of the formula[4]. Each position $p \in dom(t)$ is labeled by one points-to atomic proposition from Σ, and each atomic proposition from Σ is found in the tree, i.e. $AP(\Sigma) = \bigcup_{p \in dom(t)} t(p)$. Formally, for all positions $p \in dom(t)$, such that $t(p) \equiv x_i \mapsto (y_{i,1}, \ldots, y_{i,m_i})$, we have:

- For all $0 \leq d < \#_t(p)$, $\mathbf{alloc}(t(p.d)) = \{y_{i,j}\}$, for some $1 \leq j \leq m_i$, i.e. the children of each position correspond to points-to formulae that allocate variables pointed to by the proposition of that position.

[4] It must be possible to build the spanning tree from any free variable $x \in FV(\Sigma)$ unless the system is disconnected (in which case an error is announced and the computation aborted).

Algorithm 1 Reduction to one points-to proposition per rule

input An inductive system $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$
output An inductive system Q with one points-to proposition per rule

1: **function** SPLITSYSTEM($\mathcal{P}$)
2: $Q \leftarrow \emptyset$
3: **for all** $i = 1, \ldots, n$ **do** ▷ iterate over all predicates
4: $\overline{P}_i \leftarrow$ **empty_predicate**
5: **for all** $j = 1, \ldots, m_i$ **do** ▷ iterate over all rules of P_i
6: **assume** $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \Sigma * P_{i_1}(\mathbf{y}_1) * \ldots * P_{i_k}(\mathbf{y}_k)$
7: **choose** $root \in \mathbf{x} \cap FV(\Sigma)$
8: $t \leftarrow$ DEPTHFIRSTTRAVERSE($\Sigma, root$)
9: **if** $dom(t) = \emptyset$ **then**
10: **error**("empty rule $R_{i,j}$")
11: **if** $\bigcup_{p \in dom(t)} t(p) \neq AP(\Sigma)$ **then**
12: **error**("disconnected rule $R_{i,j}$")
13: **for all** $s = 1, \ldots, k$ **do**
14: $X_s \leftarrow \{p \in dom(t) \mid t(p) \equiv y \mapsto (z_1, \ldots, z_\ell),\ \mathbf{y}_s \cap \{z_1, \ldots, z_\ell\} \neq \emptyset\}$
15: **if** $X_s = \emptyset$ **then**
16: **error**("disconnected rule $R_{i,j}$")
17: **for all** $p \in dom(t)$ **do** ▷ create fresh predicates
18: $P_{i,j}^{p} = \{R_{i,j}^{p} \equiv\ [p = \varepsilon\ ?\ \exists \mathbf{z}]\, t(p) * \bigstar_{0 \leq d < \#t(p)} P_{i,j}^{p,d}(\mathbf{x}, \mathbf{z}) * \bigstar_{\substack{s=1,\ldots,k \\ minlex(X_s)=p}} P_{i_s}(\mathbf{y}_s)\}$
19: $\overline{P}_i \leftarrow \overline{P}_i \mid R_{i,j}^{\varepsilon}$ ▷ create root-level predicates
20: $Q \leftarrow Q \cup \{P_{i,j}^{p} \mid p \in dom(t) \setminus \{\varepsilon\}\}$
21: $Q \leftarrow Q \cup \{\overline{P}_i\}$
22: **return** CLEANUP(Q)

– For all $0 \leq d < e < \#_t(p)$, if $\mathbf{alloc}(t(p.d)) = \{y_{i,j}\}$ and $\mathbf{alloc}(t(p.e)) = \{y_{i,k}\}$, then $j < k$, i.e. the children of each node are ordered with respect to the selectors via which they are pointed to.

The spanning tree t is used to create a set of fresh predicates $P_{i,j}^{p} \equiv R_{i,j}^{p}$, one for each position $p \in dom(t) \setminus \{\varepsilon\}$, and top rules $\overline{P}_i \equiv |_{j=1}^{m_i} R_{i,j}^{\varepsilon}$, which are the only rules in which existential quantification is allowed. For each predicate occurrence $P_{i_1}(\mathbf{y}_1), \ldots, P_{i_k}(\mathbf{y}_k)$, the sets $X_1, \ldots, X_k \subseteq dom(t)$ correspond to the positions p where the actual parameters are referred to by $t(p) \equiv y \mapsto (z_1, \ldots, z_\ell)$. We chose the lexicographically minimal position from each set $X_1, \ldots, X_k$ (line 18) to place the occurrences of $P_{i_1}, \ldots, P_{i_k}$, respectively.

Finally, the new inductive system Q is cleaned (line 22) by removing (i) all unused variables from the rules and from the calls to the predicates in which they are declared, and (ii) moving existential quantifiers inside the rules where they are used. For instance, in the example below, the existential quantifier $\exists z$ has been moved from R_1 (left) to R_2 (right), because z is used in the points-to formula of R_2:

$$\begin{bmatrix} R_1(x) \equiv \exists y, z \,.\, x \mapsto y * R_2(y,z) \\ R_2(y,z) \equiv y \mapsto z \end{bmatrix} \xrightarrow{\text{CLEANUP}} \begin{bmatrix} R_1(x) \equiv \exists y \,.\, x \mapsto y * R_2(y) \\ R_2(y) \equiv \exists z \,.\, y \mapsto z \end{bmatrix}$$

The elimination of useless variables is done in reversed topological order, always processing a predicate P_i before P_j only if P_i occurs in a rule of P_j, whereas the elimination of existential quantifiers is performed in topological order, i.e. we process P_i before P_j only if P_j occurs in a rule of P_i.

It can be easily checked that DEPTHFIRSTTRAVERSE$(\Sigma, root)$ takes time $O(|\Sigma|)$, and CLEANUP(Q) takes time $O(|Q|)$. Moreover, since the inductive system Q is obtained (line 21) in time $O(|\mathcal{P}|)$, it must be the case that $|Q| = O(|\mathcal{P}|)$. Thus, the entire Algorithm 1 takes time $O(|\mathcal{P}|)$. It is not hard to check that the result of Algorithm 1 is an inductive system which is equivalent to the input, i.e. $\langle \mathcal{P}, P_i \rangle$ and $\langle \text{SPLITSYSTEM}(\mathcal{P}), \overline{P}_i \rangle$ are equivalent, for all $i = 1, \ldots, n$.

Parameter elimination. The final pre-processing step, before conversion of a rooted inductive system $\langle \mathcal{P}, P_i \rangle$ to a tree automaton, is the specialisation of $\mathcal{P}$ with respect to the predicate $P_i(x_{i,1}, \ldots, x_{i,n_i})$, and a tuple of actual parameters $\overline{\alpha} = \langle \alpha_1, \ldots, \alpha_{n_i} \rangle$, not occurring in the system. We say that a formal parameter $x_{i,k}$ of a rule $R_{i,j}(x_{i,1}, \ldots, x_{i,n_i}) \equiv \exists \mathbf{z} \,.\, \Sigma * P_{i_1}(\mathbf{y}_1) * \ldots * P_{i_m}(\mathbf{y}_m) \wedge \Pi$ is *directly propagated* to some (unique) parameter of a predicate occurrence P_{i_j}, for some $1 \leq j \leq m$, if and only if $x_{i,k} \notin FV(\Sigma)$ and $x_{i,k} \equiv (\mathbf{y}_{i_j})_\ell$, for some $0 \leq \ell < |\mathbf{y}_{i_j}|$, i.e. $x_{i,k}$ is neither allocated nor pointed to by the head of the rule, before being passed on to P_{i_j}. We denote direct propagation of parameters by the relation $x_{i,k} \rightsquigarrow x_{i_j,\ell}$, where $x_{i_j,\ell}$ is the formal parameter of P_{i_j} which is mapped to the occurrence of $(\mathbf{y}_{i_j})_\ell$. We say that $x_{i,k}$ is *propagated* to $x_{r,s}$ if $X_{i,k} \rightsquigarrow^* x_{r,s}$, where $\rightsquigarrow^*$ denotes the reflexive and transitive closure of the $\rightsquigarrow$ relation. Algorithm 2 describes the elimination of all variables that are propagated from the formal parameter tuple of a given predicate. Notice that all eliminated variables are collected in a global set `Parameters`, which will be used later, by Algorithm 4 to produce a tree automaton.

The tracking/elimination of a given formal parameter is implemented by a recursive function TRACKELIMINATE$(P_r, x_r, \texttt{del})$, where P_r is the current predicate, $x_{r,s}$ is the currently tracked formal parameter of P_r (i.e. $\texttt{Tracked} \rightsquigarrow^* x_{r,s}$ is an invariant for every call to TRACKELIMINATE), and $\texttt{del} = \textbf{true}$ if and only if $x_{r,s}$ is to be removed from the definition of P_r. If the currently tracked parameter $x_{r,s}$ is either allocated in a rule of P_r or it is not propagated further, then every occurrence of $x_{r,s}$ is replaced with `Tracked` (line 8). Otherwise, if the parameter is propagated to P_j as $x_{j,\ell}$ (line 10) and it is referenced by the current rule (line 13), then it will not be removed any longer from the system (line 15). In this case, we keep tracking it only to place the global variable `Tracked` in the right place (line 8). Otherwise, if the parameter is not referenced, it will be removed completely from the rule (line 18). Finally, the old predicates in $Q \cap \mathcal{P}$, which have become unreachable from the P_i, due to the insertion of the new ones (Q_r), are removed from the system by a call to the TRIM procedure. The running time of the algorithm is linear in the size of $\mathcal{P}$ (i.e. each rule is visited at most once) and the size of the output system is reduced with respect to the input, i.e. $|Q| \leq |\mathcal{P}|$.

Example 6 (cont. of Ex. 1). As an example of parameter elimination, let us take the predicate `DLL` introduced in Sect. 2.4, called as $\texttt{DLL}(\texttt{a}, \texttt{b}, \texttt{c}, \texttt{d})$. After the parameter elimination and renaming the newly created predicates, we have a call Q_1 (without parameters) of the following inductive system:

$$Q_1() \equiv \texttt{a} \mapsto (\texttt{d}, \texttt{b}) \ \wedge \ \texttt{a} = \texttt{c} \mid \exists x.\ \texttt{a} \mapsto (x, \texttt{b}) * Q_2(x, \texttt{a})$$
$$Q_2(hd, p) \equiv hd \mapsto (\texttt{d}, p) \ \wedge \ hd = \texttt{c} \mid \exists x.\ hd \mapsto (x, p) * Q_2(x, hd)$$

4.2 Conversion to tree automata

The following algorithm takes as input a rooted inductive system $\langle \mathcal{P}, P_i \rangle$ such that P_i has no formal parameters, the head of each rule in $\mathcal{P}$ is of the form $\exists \mathbf{z} \,.\, \alpha \mapsto (y_0, \ldots, y_{m-1}) \wedge \Pi$, and for each $x \in FV(\Pi)$, we have (i) $\alpha \in [x]_\Pi$ and (ii) $[x]_\Pi \cap \{y_0, \ldots, y_{m-1}\} = \emptyset$. It is easy to see that these conditions are ensured by the first two steps of pre-processing (i.e. the elimination of equalities and the reduction to one points-to proposition per rule). The first step of the conversion computes several sets of parameters that are called *signatures* and defined as follows:

$$\begin{aligned}
\mathtt{sig}_j^{fw} &= \{x_{j,k} \mid x_{j,k} \text{ is allocated in every rule of } P_j\} \cap \\
&\quad \{x_{j,k} \mid \text{there is an edge to } (\mathbf{y})_k \text{ in every occurrence } P_j(\mathbf{y}) \text{ of } \mathcal{P}\} \\
\mathtt{sig}_j^{bw} &= \{x_{j,k} \mid \text{there is an edge to } x_{j,k} \text{ in every rule of } P_j\} \cap \\
&\quad \{x_{j,k} \mid (\mathbf{y})_k \text{ is allocated for every occurrence } P_j(\mathbf{y}) \text{ of } \mathcal{P}\} \\
\mathtt{sig}_j^{eq} &= \{x_{j,1}, \ldots, x_{j,n_j}\} \setminus (\mathtt{sig}_j^{fw} \cup \mathtt{sig}_j^{bw})
\end{aligned}$$

where $\langle x_{j,1}, \ldots, x_{j,n_j} \rangle$ is the tuple of formal parameters of the predicate P_j, in $\mathcal{P}$, for all $j = 1, \ldots, n$. The signatures of the system are computed via Algorithm 3 .

The result of the algorithm can be used to implement a sufficient test of locality: a given system is local if $sig_i^{eq} = \emptyset$, for each i. This simple test turned out to be powerful enough for all the examples that we considered in our experiments in Section 6.

Example 7 (cont. of Ex. 6). The signatures for the system in Example 6 are:

- $sig_1^{fw} = sig_1^{bw} = sig_1^{eq} = \emptyset$
- $sig_2^{fw} = \{0\}, sig_2^{bw} = \{1\}, sig_2^{eq} = \emptyset$

The fact that, for each i, we have $sig_i^{eq} = \emptyset$ informs us that the system is local.

Next, we define a normal form for the quasi-canonical tiles, used as alphabet symbols in the rules of the tree automaton. Given a quasi-canonical tile $T = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{k-1} \rangle$, for some $k \geq 0$, where $\varphi \equiv (\exists z)\ z \mapsto (y_0, \ldots, y_{m-1}) \wedge \Pi$, for some $m > 0$, let $\xi_0, \ldots, \xi_m$ be unique variable namesi. Let σ be the substitution defined as $\sigma(z) = \xi_0$, and $\sigma(y_i) = \xi_{i+1}$, for all $0 \leq i < m$. For all variables $\mathbf{x}_{-1} \cup \mathbf{x}_0 \cup \ldots \cup \mathbf{x}_{k-1} = \{\alpha_1, \ldots, \alpha_n\}$, we chose unique names $\xi_{m+1}, \ldots, \xi_{m+n}$, according to the order in which they appear in the ports $\mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{k-1}$, respectively, and we extend the substitution further by defining $\sigma(\alpha_i) = \xi_{m+i}$, for all $i = 1, \ldots, n$. We define the *normal tile* $\overline{T} = \langle \varphi[\sigma], \mathbf{x}_{-1}[\sigma], \mathbf{x}_0[\sigma], \ldots, \mathbf{x}_{k-1}[\sigma] \rangle$. For a quasi-canonical tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$, we define the *normal tree* $\bar{t}$ as $\bar{t}(p) = \overline{t(p)}$, for all $p \in dom(t)$.

The SL2TA function (Algorithm 4) builds a tree automaton A, for a rooted system $\langle \mathcal{P}, P_i \rangle$. For each rule in the system, the algorithm creates a quasi-canonical tile, where the input and output ports $\mathbf{x}_i$ are factorized as $\mathbf{x}_i^{fw} \cdot \mathbf{x}_i^{bw} \cdot \mathbf{x}_i^{eq}$, according to the precomputed signatures. The backward part of the input port $\mathbf{x}_{-1}^{bw}$ and the forward parts of the output ports $\mathbf{x}_i^{fw}$, for $i \geq 0$, are sorted according to the order of incoming selector edges from the single points-to formula $\alpha \mapsto (\overline{\beta})$, by the SELECTORSORT function (lines 18 and 28). The output ports x_i, $i \geq 0$, are sorted within the tile, according to the order of the selector edges pointing to $(\mathbf{x}_i^{fw})_0$, for each $i = 0, 1, \ldots$ (function SORTTILE, line 31). Finally, each predicate name P_j is associated a state q_j (line 33), and for each inductive rule, the algorithm creates a transition rule in the tree automaton (line 34). The final state corresponds to the root of the system (line 35). The following lemma summarizes the tree automata construction.

Lemma 6. *Given a rooted inductive system* $\langle \mathcal{P}, P_i(\mathbf{x})\rangle$ *where* $\mathcal{P} = \left\{P_i \equiv |_{j=1}^{m_i} R_{i,j}\right\}_{i=1}^{n}$ *and* $\mathbf{x} = \langle x_{i,1}, \ldots, x_{i,n_i}\rangle$*, and a vector* $\overline{\alpha} = \langle \alpha_1, \ldots, \alpha_{n_i}\rangle$ *of variables not used in* $\mathcal{P}$*. Then, for every state* S*, we have* $S \models P_i(\overline{\alpha})$ *if and only if there exists* $t \in \mathcal{L}(A)$ *such that* $S \models \Phi(t)$ *where* $A = \text{SL2TA}(\mathcal{P}, i, \overline{\alpha})$*. Moreover,* $|A| = O(|\mathcal{P}|)$*.*

Example 8 (cont. of Ex. 7). The automaton corresponding to the DLL system called by $\texttt{DLL}(\texttt{a},\texttt{b},\texttt{c},\texttt{d})$ is $A = \langle \Sigma, \{q_1, q_2\}, \Delta, \{q_1\}\rangle$ where:

$$\Delta = \left\{ \begin{array}{ll} \overline{\langle \texttt{a} \mapsto (\texttt{b},\texttt{d}) \wedge \texttt{a} = \texttt{c}, \emptyset\rangle}() \to q_1 \quad \overline{\langle \texttt{a} \mapsto (x,\texttt{b}), \emptyset, (x,\texttt{a})\rangle}(q_2) & \to q_1 \\ \overline{\langle \exists hd'.hd' \mapsto (\texttt{d}, p) \wedge hd = \texttt{c} \wedge hd' = hd, (hd, p)\rangle}() & \to q_2 \\ \overline{\langle \exists hd'.hd' \mapsto (x, p) \wedge hd' = hd, (hd, p), (x, hd)\rangle}(q_2) & \to q_2 \end{array} \right\}$$

4.3 Rotation of Tree Automata

This section describes the algorithm that produces the closure of a quasi-canonical tree automaton (i.e. a tree automaton recognizing quasi-canonical trees only) under rotation. The result of the algorithm is a tree automaton A^r that recognizes all trees $u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$ such that $t \sim^{qc} u$, for some tree t recognized by $A = \langle Q, \Sigma, \Delta, F\rangle$. Algorithm 5 describes the rotation closure. The result of Algorithm 5 (function SL2TA) is a language-theoretic union of A and automata A_ρ, one for each rule ρ of A. The idea behind the construction of $A_\rho = \langle Q_\rho, \Sigma, \Delta_\rho, \{q_\rho^f\}\rangle$ can be understood by considering a tree $t \in \mathcal{L}(A)$, a run $\pi : dom(t) \to Q$, and a position $p \in dom(t)$, which is labeled with the right hand side of the rule $\rho = T(q_1, \ldots, q_k) \to q$ of A. Then $\mathcal{L}(A_\rho)$ will contain the rotated tree u, i.e. $t \sim_r^{qc} u$, where the significant position p is mapped into the root of u by the rotation function r, i.e. $r(p) = \varepsilon$. To this end, we introduce a new rule $T_{new}(q_0, \ldots, q_j, q^{rev}, q_{j+1}, \ldots, q_n) \to q_\rho^f$, where the tile T_{new} mirrors the change in the structure of T at position p, and $q^{rev} \in Q_\rho$ is a fresh state corresponding to q. The construction of A_ρ continues recursively, by considering every rule of A that has q on the left hand side: $U(q'_1, \ldots, q, \ldots, q'_\ell) \to s$. This rule is changed by swapping the roles of q and q' and producing a rule $U_{new}(q'_1, \ldots, s^{rev}, \ldots q'_\ell) \to q^{rev}$, where U_{new} mirrors the change in the structure of U. Intuitively, the states $\{q^{rev} | q \in Q\}$ mark the unique path from the root of u to $r(\varepsilon) \in dom(u)$. The recursion stops when either (i) s is a final state of A, (ii) the tile U does not specify a forward edge in the direction marked by q, or (iii) all states of A have been visited.

The following lemma proves the correctness of Algorithm 5.

Lemma 7. *Let* $A = \langle Q, \mathcal{T}^{qc}, \Delta, F\rangle$ *be a tree automaton. Then* $\mathcal{L}(A^r) = \{\overline{u} \mid u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc},\ \exists t \in \mathcal{L}(A)\ .\ u \sim^{qc} t\}$*. Moreover, the size of* A^r *is of the order of* $O(|A|^2)$*.*

Proof. "$\subseteq$" Let $u \in \mathcal{L}(A^r)$ be a normal quasi-canonical tree. Then $u \in \mathcal{L}(A) \cup \bigcup_{\rho \in \Delta} \mathcal{L}(A_\rho)$ (lines 1 and 14 in Alg. 5). If $u \in \mathcal{L}(A)$, then we choose $t = u$ and trivially $t \sim^{qc} u$. Otherwise $u \in \mathcal{L}(A_\rho)$, where $A_\rho = \langle Q_\rho, \mathcal{T}^{qc}, \Delta_\rho, \{q_\rho^f\}\rangle$, for some $\rho \in \Delta$ (line 13). Let $\pi : dom(u) \to Q_\rho$ be the run of A_ρ on u. Also let $p_0 \in dom(u)$ be the maximal (i.e. of maximal length) position such that $u(p_0) = q^{rev}$, for some $q^{rev} \in Q^{rev}$ (line 7). Let $u_{|_{p_0.i}}$ be the subtrees of u rooted at the children of p_0, for all $i = 0, \ldots, \#_u(p_0) - 1$. Since p_0 is the maximal position of π to be labeled by some state in Q^{rev}, it is easy to see that $\pi_{|_{p_0.i}}$

are labeled by states in Q, hence $\pi_{|_{p_0.i}}$ are runs of A over $t_{|_{p_0.i}}$, for all $i = 0,\ldots,\#_u(p_0)-1$ (the only rules of A_ρ involving only states from Q are the rules of A, cf. line 8).

We build a quasi-canonical tree $t : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$ and a run $\theta : dom(t) \to Q$ of A on t top-down as follows. Let $p_0, p_1, \ldots, p_n = \varepsilon$ be the path composed of the prefixes of p_0, i.e. $p_i = p_{i-1}.(-1)$, for all $i = 1,\ldots,|p_0| = n$. We build t, θ, and a path $\overline{p_0} = \varepsilon, \overline{p_1} = j_0, \ldots, \overline{p_n} \in dom(t)$, by induction on this path.

For the base case, let $u(p_0) = \langle \varphi, \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{eq}, \mathbf{x}_0, \ldots, \mathbf{x}_{\#_u(p)-1} \rangle$ be the tile which labels u at position p_0. Let $t(\varepsilon) = \langle \varphi, \emptyset, \mathbf{x}_0, \ldots, \mathbf{x}_{j_0-1}, \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{eq}, \mathbf{x}_{j_0}, \ldots, \mathbf{x}_{\#_u(p)-1} \rangle$, where j_0+1 is the canonical position of the output port $\mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \mathbf{x}_{-1}^{eq}$, in $t(\varepsilon)$, according to the order of selector edges in φ. Then A_ρ has a rule:

$$u(p_0)(\pi(p_0.0), \ldots, \pi(p_0.(\#_u(p)-1))) \to \pi(p)$$

such that $\{\pi(p_0.0), \ldots, \pi(p_0.(\#_u(p)-1)\} \cap Q^{rev} = \emptyset$. Since p_0 is the maximal position labeled with a state from Q^{rev}, this rule was generated at line 9 in Alg. 5. It follows that A has a rule

$$t(\varepsilon)(\pi(p_0.0), \ldots, \pi(p_0.(\#_u(p_0)-1))) \to q$$

for some final state $q \in F$ (line 7). Then let $\theta(\varepsilon) = q$. We further define:

$$t_{|_i} = \begin{cases} u_{|_{p_0.i}} & \text{if } i = 0,\ldots,j_0-1 \\ u_{|_{p_0.(i+1)}} & \text{if } i = j_0+1,\ldots,\#_u(p_0)-1 \end{cases} \qquad \theta_{|_i} = \begin{cases} \pi_{|_{p_0.i}} & \text{if } i = 0,\ldots,j_0-1 \\ \pi_{|_{p_0.(i+1)}} & \text{if } i = j_0+1,\ldots,\#_u(p_0)-1 \end{cases}$$

For the induction step, for each $0 < i \leq n$, we have $p_{i-1} = p_i.k_i$, for some $0 \leq k_i < \#_u(p_{i-1})$. We have $u(p_i) = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_{\#_u(p_i)-1} \rangle$, and define:

$$t(\overline{p_i}) = \langle \varphi,\ \mathbf{x}_{k_i}^{bw} \cdot \mathbf{x}_{k_i}^{fw} \cdot \mathbf{x}_{k_i}^{eq},\ \mathbf{x}_0,\ \ldots,\ \mathbf{x}_{j_i-1},\ \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{eq},\ \mathbf{x}_{j_i},\ \ldots,\ \mathbf{x}_{\#_u(p_i)-1} \rangle$$

where j_i+1 is the canonical position of the port $\mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{eq}$ given by the the selector edges in φ. Moreover, A_ρ has a rule:

$$u(p_i)(\pi(p_i.0), \ldots, \pi(p_i.(\#_u(p_i)-1))) \to \pi(p_i)$$

where $\pi(p_i.k_i) = \pi(p_{i-1}) \in Q^{rev}$, $\pi(p_i) \in Q^{rev}$ if $0 \leq i < n$, and $\pi(p_n) = q_\rho^f$. Moreover, this rule was introduced at line 18, if $i < n$, or at line 11, if $i = n$. Let $i < n$ (the case $i = n$ uses a similar argument). If $\pi(p_{i-1}) = s^{rev}$ and $\pi(p_i) = q^{rev}$, then A must have a rule:

$$t(\overline{p_i})(\pi(p_i.0), \ldots, \pi(p_i.(j_i-1)), q, \pi(p_i.j_i), \ldots, \pi(p_i.(\#_u(p_i)-1))) \to s$$

Let $\overline{p_{i+1}} = \overline{p_i}.j_i$ and $\theta(\overline{p_{i+1}}) = q$. We further define:

$$t_{|_{\overline{p_i}.\ell}} = \begin{cases} u_{|_{p_i.\ell}} & \text{if } \ell = 0,\ldots,j_i-1 \\ u_{|_{p_i.(\ell+1)}} & \text{if } \ell = j_i+1,\ldots,\#_u(p_i)-1 \end{cases} \qquad \theta_{|_{\overline{p_i}.\ell}} = \begin{cases} \pi_{|_{p_i.\ell}} & \text{if } \ell = 0,\ldots,j_i-1 \\ \pi_{|_{p_i.(\ell+1)}} & \text{if } i = j_i+1,\ldots,\#_u(p_i)-1 \end{cases}$$

We define the following rotation function $r : dom(t) \to dom(u)$. For each $i = 0,\ldots,n$, we have $r(\overline{p_i}) = p_i$, and:

$$r(\overline{p_i}.\ell) = \begin{cases} r(p_i.\ell) & \text{if } \ell = 0,\ldots,j_i-1 \\ r(p_i.(\ell+1)) & \text{if } \ell = j_i,\ldots,\#_u(p_i)-1 \end{cases}$$

It is easy to check that indeed $t \sim_r^{qc} u$, and that θ is a run of A over t.
"$\supseteq$" Let $u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc}$ be a quasi-canonical tree such that $t \sim_r^{qc} u$, for some $t \in \mathcal{L}(A)$ and a bijective function $r : dom(t) \to dom(u)$. Let $\pi : dom(t) \to Q$ be a run of A over t. We build an accepting run θ of A^r over the normal tree $\overline{u}$. Let $p_0 \in dom(t)$ such that $r(p_0) = \varepsilon$ is the root of u. If $p_0 = \varepsilon$, it is easy to show that $dom(u) = dom(t)$ and $r(p) = p$, for all $p \in dom(t)$, because both t and u are quasi-canonical trees, and the order of children is given by the order of selector edges in the tiles labeling the trees. Moreover, the normal form of these tiles is identical, i.e. $\overline{u}(p) = t(p)$, for all $p \in dom(t)$, hence $u \in \mathcal{L}(A) \subseteq \mathcal{L}(A)^r$ (cf. line 1).

Otherise, if $p_0 \neq \varepsilon$, we consider the sequence of prefixes of p_0, defined as $p_i = p_{i-1}.(-1)$, for all $i = 1, \ldots, n = |p_0|$. Applying Lemma 1 to $p_0, \ldots, p_n$ successively, we obtain a path $\varepsilon = r(p_0), r(p_1), \ldots, r(p_n) \in dom(u)$, such that $r(p_{i+1}) = r(p_i).d_i$, for all $i = 0, \ldots, n-1$, and some positive directions $d_0, \ldots, d_{n-1} \in \mathcal{D}_+(u)$.

We build the run θ by induction on this path. For the base case, let $t(p_0) = q_0$. Then A has a rule $\rho = (t(p_0)(q_1, \ldots, q_{\#_t(p_0)}) \to q_0) \in \Delta$. By the construction of A^r, cf. line 11, A_ρ has a rule

$$(\overline{u}(\varepsilon))(q_1, \ldots, q_{d_0}, q_0^{rev}, q_{d_0+1}, \ldots, q_{\#_t(p_0)}) \to q_\rho^f.$$

We define $\theta(\varepsilon) = q_\rho^f$ and:

$$\theta_{|_i} = \begin{cases} \pi_{|_{p_0.i}} & \text{if } i = 0, \ldots, d_0 - 1 \\ \pi_{|_{p_0.(i+1)}} & \text{if } i = d_0 + 1, \ldots, \#_t(p_0) - 1 \end{cases}$$

For the induction step, we denote $q_i = \pi(p_i)$, for all $0 < i \leq n$. Then A has a rule

$$(t(p_i))(\pi(i.0), \ldots, \pi(i.(k_i - 1)), q_i, \pi(i.(k_i + 1)), \ldots, \pi(i.(\#_t(p_i) - 1))) \to q_{i+1},$$

for each $0 < i \leq n$. By the construction of A^r, cf. line 18, A_ρ has a rule:

$$(\overline{u(r(p_i))})(\pi(i.0), \ldots, \pi(i.(d_i - 1)),\ q_{i+1}^{rev},\ \pi(i.(d_i + 1)), \ldots, \pi(i.(\#_t(p_i) - 1))) \to q_i^{rev}.$$

We define $\theta(r(p_i)) = q_i^{rev}$ and:

$$\theta_{|_{r(p_i).\ell}} = \begin{cases} \pi_{|_{r(p_i).\ell}} & \text{if } \ell = 0, \ldots, d_i - 1 \\ \pi_{|_{r(p_i).(\ell+1)}} & \text{if } i = d_i + 1, \ldots, \#_t(p_i) - 1 \end{cases}$$

It is not difficult to prove that θ is a run of A_ρ, and, moreover, since $\theta(\varepsilon) = q_\rho^f$, it is an accepting run, hence $\overline{u} \in \mathcal{L}(A_\rho) \subseteq A^r$ (cf. line 14).

Concerning the size of A^r, notice that $|A^r| \leq |A| + \sum_{\rho \in \Delta} |A_\rho|$ where $A_\rho = \langle Q_\rho, \Sigma, \Delta_\rho, \{q_\rho^f\} \rangle$. We have that $\|\Delta_\rho\| \leq \|\Delta\|$, and for each rule $\tau \in \Delta_\rho \setminus \Delta$ created at lines 9, 11, or 18, there exists a rule $\nu \in \Delta$ such that $|\tau| \leq |\nu| + 1$. Moreover, $\|\Delta_\rho \setminus \Delta\| \leq \|Q\|$ since we introduce a new rule for each state in the set `visited` $\subseteq Q$. Hence $|A_\rho| = |A| + \sum_{\tau \in \Delta_\rho \setminus \Delta} |\tau| \leq |A| + |A| + \|Q\| \leq 3|A|$. Hence $|A^r| \leq 3|A|^2 = O(|A|^2)$. $\square$

The main result of this paper is given by the following theorem. The entailment problem for inductive systems is reduced, in polynomial time, to a language inclusion problem for tree automata. The inclusion test is sound (if the answer is yes, the entailment holds), and moreover, complete assuming that both systems are local.

Theorem 1. *Let $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ be a system of inductive definitions. Then, for any two predicates $P_i(x_{i,1},\ldots,x_{i,n_i})$ and $P_j(x_{j,1},\ldots,x_{j,n_j})$ of $\mathcal{P}$ such that $n_i = n_j$ and for any variables $\overline{\alpha} = \{\alpha_1,\ldots,\alpha_{n_i}\}$ not used in $\mathcal{P}$, the following holds:*

- *[Soundness] $P_i(\overline{\alpha}) \models_{\mathcal{P}} P_j(\overline{\alpha})$ if $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2^r)$,*
- *[Completness] $P_i(\overline{\alpha}) \models_{\mathcal{P}} P_j(\overline{\alpha})$ only if $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2^r)$ provided that the rooted systems $\langle \mathcal{P}, P_i\rangle$ and $\langle \mathcal{P}, P_j\rangle$ are both local*

where $A_1 = \text{SL2TA}(\mathcal{P}, i, \overline{\alpha})$ and $A_2 = \text{SL2TA}(\mathcal{P}, j, \overline{\alpha})$.

Proof (Soundness). Let S be a state such that $S \models P_i(\overline{\alpha})$. By Lemma 6, there exists a tree $t \in \mathcal{L}(A_1)$ such that $S \models \Phi(t)$. Since $t \in \mathcal{L}(A_1) \subseteq \mathcal{L}(A_2^r)$, by Lemma 7, there exists a quasi-canonical tree $u \in \mathcal{L}(A_2)$ such that $t \sim^{qc} u$. By Lemma 5, we have $S \models \Phi(u)$, and by Lemma 6 again, we obtain $S \models P_j(\overline{\alpha})$. Hence $P_i(\overline{\alpha}) \models_{\mathcal{P}} P_j(\overline{\alpha})$.

[Completness] If $\langle \mathcal{P}, P_i\rangle$ and $\langle \mathcal{P}, P_j\rangle$ are local rooted systems, it follows that A_1 and A_2 recognize only normal canonical trees. Let $t \in \mathcal{L}(A_1)$ be a normal canonical tree. It is not difficult to see that starting from t, one can build a state S and its spanning tree such that the right-hand side of Lemma 3 holds. Then, by Lemma 3, one gets that $S \models \Phi(t)$. By Lemma 6, we have $S \models P_i(\overline{\alpha})$, and since $P_i(\overline{\alpha}) \models_{\mathcal{P}} P_j(\overline{\alpha})$, we have $S \models P_j(\overline{\alpha})$. By Lemma 6 again, we obtain a canonical tree $u \in \mathcal{L}(A_2)$ such that $S \models \Phi(u)$. Since $S \models \Phi(t)$ and $S \models \Phi(u)$, by Lemma 4, we have $t \sim^{c} u$, and since t is normal, we obtain $t \in \mathcal{L}(A_2^r)$, by Lemma 7. Hence $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2^r)$. □

Example 9 (cont. of Ex. 8). The automaton A^r corresponding to the rotation closure of the automaton A is $A^r = \langle \Sigma, \{q_1, q_2, q_2^{rev}, q_{fin}\}, \Delta, \{q_1, q_{fin}\}\rangle$ where:

$$\Delta = \left\{ \begin{array}{ll}
\overline{\langle \mathtt{a} \mapsto (\mathtt{b},\mathtt{d}) \wedge \mathtt{a} = \mathtt{c}, \emptyset\rangle}() \to q_1 \quad \overline{\langle \mathtt{a} \mapsto (x,\mathtt{b}), \emptyset, (x,\mathtt{a})\rangle}(q_2) & \to q_1 \\
\overline{\langle \exists hd'.hd' \mapsto (\mathtt{d},p) \wedge hd = \mathtt{c} \wedge hd' = hd, (hd,p)\rangle}() & \to q_2 \\
\overline{\langle \exists hd'.hd' \mapsto (x,p) \wedge hd' = hd, (hd,p), (x,hd)\rangle}(q_2) & \to q_2 \\
\overline{\langle \exists hd'.hd' \mapsto (d,p) \wedge hd = c \wedge hd' = hd, \emptyset, (p,hd)\rangle}(q_2^{rev}) & \to q_{fin} \\
\overline{\langle \mathtt{a} \mapsto (x,\mathtt{b}), (\mathtt{a},x)\rangle}() & \to q_2^{rev} \\
\overline{\langle \exists hd'.hd' \mapsto (x,p) \wedge hd' = hd, (hd,x), (p,hd)\rangle}(q_2^{rev}) & \to q_2^{rev} \\
\overline{\langle \exists hd'.hd' \mapsto (x,p) \wedge hd' = hd, \emptyset, (x,hd), (p,hd)\rangle}(q_2, q_2^{rev}) & \to q_{fin}
\end{array} \right\}$$

5 Complexity

In this section, we prove tight complexity bounds for the entailment problem in the fragment of SL with inductive definitions under consideration (i.e. with the restrictions defined in Sec. 2.4). The first result shows the need for *connected* rules in the system as allowing unconnected rules leads to the undecidability of the entailment problem. As a remark, the general undecidability of entailments for SL with inductive definitions has already been proved in [1]. In addition to their result, our proof stresses the fact that undecidability occurs due the lack of connectivity within some rules.

Theorem 2. *The entailment problem $P_i \models_{\mathcal{P}} P_j$, for some $1 \leq i, j \leq n$, is undecidable for inductive systems $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$ that can have unconnected rules.*

Proof. A context-free grammar is a tuple $G = \langle \mathcal{X}, \Sigma, \delta \rangle$ where $\mathcal{X}$ is a finite nonempty set of *nonterminals*, Σ is a finite nonempty *alphabet* such that $\mathcal{X} \cap \Sigma = \emptyset$, and $\delta \subseteq \mathcal{X} \times (\Sigma \cup \mathcal{X})^*$ is a finite set of *productions*. Given two strings $u, v \in (\Sigma \cup \mathcal{X})^*$, we define a *step* $u \Longrightarrow_G v$ if there exists a production $(X, w) \in \delta$ and some words $y, z \in (\Sigma \cup \mathcal{X})^*$ such that $u = yXz$ and $v = ywz$. We denote by $\Longrightarrow_G^*$ the reflexive and transitive closure of the $\Longrightarrow_G$ relation. The *language* of a grammar G generated by a nonterminal $X \in \mathcal{X}$ is defined as $L_X(G) = \{w \in \Sigma^* \mid X \Longrightarrow_G^* w\}$. It is known that the inclusion problem $L_X(G) \subseteq L_Y(G)$, for some nonterminals $X, Y \in \mathcal{X}$, is undecidable as originally proved in [2]. We reduce from this problem to entailment within unconnected inductive systems.

Let $\Sigma = \{\sigma_1, \ldots, \sigma_N\}$ be the alphabet of G. We define the set of selectors

$$Sel = \{1, \ldots, \lceil \log_2 N \rceil + 1\}$$

and, for each alphabet symbol σ_K, we define a basic SL formula $\varphi_K(x, y) \equiv x \mapsto (\beta, y)$ where, for all $0 \leq i < \lceil \log_2 N \rceil$:

$$\beta_i = \begin{cases} x & \text{if 1 occurs on position } i \text{ in the binary encoding of } K \\ \textbf{nil} & \text{if 0 occurs on position } i \text{ in the binary encoding of } K \end{cases}$$

A word $w = \sigma_{i_1} \cdot \ldots \cdot \sigma_{i_n} \in \Sigma^*$ is encoded by the formula

$$\varphi_w(x, y) \equiv \exists x_1 \ldots \exists x_{n-1} \,.\, \varphi_{i_1}(x, x_1) * \ldots \varphi_{i_n}(x_{n-1}, y).$$

We define a predicate name $P_X(x, y)$ for each nonterminal $X \in \mathcal{X}$, and for each production $\pi \equiv (X, w_1 \cdot X_{i_1} \cdot \ldots \cdot w_n \cdot X_{i_n} \cdot w_{n+1}) \in \delta$ where $w_1, \ldots, w_n \in \Sigma^*$ are words and $X_{i_1}, \ldots, X_{i_n} \in \mathcal{X}$ are non-terminals, we have the rule:

$$\begin{aligned} R_\pi(x, y) \equiv \exists x_1 \ldots \exists x_{2n} \,.\, & \varphi_{w_1}(x, x_1) * P_{X_{i_1}}(x_1, x_2) * \ldots * \\ & \varphi_{w_n}(x_{2n-2}, x_{2n-1}) * P_{X_{i_n}}(x_{2n-1}, x_{2n}) * \varphi_{w_{n+1}}(x_{2n}, y). \end{aligned}$$

Finally, each predicate is defined as $P_X \equiv \mid_{(X,w) \in \delta} R_{(X,w)}$, and the inductive system is $\mathcal{P}_G = \{P_X \mid X \in \mathcal{X}\}$. It is immediate to check that $P_X(x, y) \models_{\mathcal{P}_G} P_Y(x, y)$ if and only if $L_X(G) \subseteq L_Y(G)$. □

The second result of this section provides tight complexity bounds for the entailment problem for connected systems with the restrictions defined in Sec. 2.4. We must point out that EXPTIME-hardness of entailments in the fragment of [13] was already proved in [1]. The result below is stronger as the fragment under consideration is a restriction of the fragment from [13]. In particular, we do not allow branching parameter copying and do not consider checking entailments between formulae with empty spatial heads like, e.g., $\exists z \,.\, \texttt{DLL}(x, z) * \texttt{DLL}(z, y) \models \texttt{DLL}(x, y)$. On the other hand, these restrictions allow us to set an EXPTIME upper bound on the entailment problem for connected systems.

Theorem 3. *The entailment problem $P_i \models_{\mathcal{P}} P_j$, for some $1 \leq i, j \leq n$, is EXPTIME-complete for connected inductive systems $\mathcal{P} = \{P_i \equiv \mid_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$.*

Proof. We reduce from and to the inclusion problem for non-deterministic (bottom-up) tree automata. It is known that the language inclusion problem for non-deterministic tree automata is EXPTIME-complete (see, e.g. Corollary 1.7.9 in [8]).

For the hardness part, let $A_1 = \langle Q_1, \Sigma, \Delta_1, F_1 \rangle$ and $A_2 = \langle Q_2, \Sigma, \Delta_2, F_2 \rangle$ be two tree automata over the same alphabet $\Sigma = \{\sigma_1, \ldots, \sigma_N\}$. We assume w.l.o.g. that $Q_1 \cap Q_2 = \emptyset$. We define the set of selectors

$$Sel = \{1, \ldots, \lceil \log_2 N \rceil + \max\{\#(\sigma) \mid \sigma \in \Sigma\}\}.$$

For each alphabet symbol σ_K, $1 \leq K \leq N$ and each rule $\rho \equiv (\sigma_K(q_1, \ldots, q_n) \to q) \in \Delta_1 \cup \Delta_2$, we define a basic SL formula $\varphi^{\rho}_K(x, y_1, \ldots, y_n) \equiv x \mapsto (\beta, y_1, \ldots, y_n)$ where, for all $0 \leq i < \lceil \log_2 N \rceil$:

$$\beta_i = \begin{cases} x & \text{if 1 occurs on position } i \text{ in the binary encoding of } K \\ \textbf{nil} & \text{if 0 occurs on position } i \text{ in the binary encoding of } K \end{cases}$$

For each state $q \in Q_i$ of A_i, we consider a predicate name $P_q(x)$ and an extra predicate name $P_i(x)$, for both $i = 1, 2$. For each transition rule $\rho \equiv (\sigma_K(q_1, \ldots, q_n) \to q)$, of either A_1 or A_2, we define the following connected inductive rule:

$$R_\rho(x) \equiv \exists y_1 \ldots \exists y_n \,.\, \varphi^{\rho}_K(x, y_1, \ldots, y_n) * P_{q_1}(y_1) * \ldots * P_{q_n}(y_n)$$

Then we define the predicates:

$$\begin{array}{l} P_q = \{R_q \mid (\sigma(q_1, \ldots, q_n) \to q) \in \Delta_1 \cup \Delta_2\} \\ P_i = \{R_q \mid (\sigma(q_1, \ldots, q_n) \to q) \in \Delta_i \text{ and } q \in F_i\} \text{ for both } i = 1, 2 \end{array}$$

Finally, we have $\mathcal{P} = \{P_q \mid (\sigma(q_1, \ldots, q_n) \to q) \in \Delta_1 \cup \Delta_2\} \cup \{P_1, P_2\}$. It is easy to check that $P_1(x) \models_{\mathcal{P}} P_2(x)$ if and only if $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2)$. Hence the entailment problem is EXPTIME-hard.

For the EXPTIME-completness part, Theorem 1 shows that any entailment problem $P_i \models_{\mathcal{P}} P_j$, for a connected system $\mathcal{P}$, can be reduced, in polynomial time, to a language inclusion problem between tree automata $\mathcal{L}(A_{P_i}) \subseteq \mathcal{L}(A^r_{P_j})$. Hence the entailment problem is in EXPTIME. □

6 Experiments

We implemented a prototype tool [12] that takes as input two rooted inductive systems $\langle \mathcal{P}_{lhs}, P_{lhs} \rangle$ and $\langle \mathcal{P}_{rhs}, P_{rhs} \rangle$ with eliminated parameters (i.e. the result of Alg. 2). Table 1 lists the entailment queries on which we tried out our tool. The upper part of the table contains local inductive systems whereas the bottom part contains non-local systems. Apart from the DLL and TLL predicates from Sect. 2.4, the considered entailment queries contain the following predicates : $\texttt{DLL}_{rev}$ that encodes a DLL from the end, $\texttt{TREE}_{pp}$ encoding trees with parent pointers, $\texttt{TREE}^{rev}_{pp}$ that encodes trees with parent pointers starting from an arbitrary leaf, $\texttt{TLL}_{pp}$ encoding TLLs with parent pointers, and $\texttt{TLL}^{rev}_{pp}$ which encodes TLLs with parent pointers starting from their left most leaf. These predicates are defined in Fig. 6. Columns $|A_{lhs}|$, $|A_{rhs}|$, and $|A^r_{rhs}|$ of Table 1 provide information about the number of states and transitions of the respective automata. The tool answered all queries correctly (despite the incompleteness for non-local systems), and the running times were all under 1 sec. on a standard PC (Intel Core2 CPU, 3GHz, 4GB RAM).

Table 1. Experimental results. The upper table contains local systems, while the lower table non-local ones. Sizes of initial TA (col. 3,4) and rotated TA (col. 5) are in numbers of states/transitions

Entailment $LHS \models RHS$	Answer	$\|A_{lhs}\|$	$\|A_{rhs}\|$	$\|A^r_{rhs}\|$
$\mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{DLL}_{rev}(a,\mathbf{nil},c,\mathbf{nil})$	True	2/4	2/4	5/8
$\mathtt{DLL}_{rev}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil})$	True	2/4	2/4	5/8
$\mathtt{DLL}_{mid}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil})$	True	4/8	2/4	5/8
$\mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{DLL}_{mid}(a,\mathbf{nil},c,\mathbf{nil})$	True	2/4	4/8	12/18
$\exists x,n,b.\ x \mapsto (n,b) * \mathtt{DLL}_{rev}(a,\mathbf{nil},b,x) * \mathtt{DLL}(n,x,c,\mathbf{nil}) \models \mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil})$	True	3/5	2/4	5/8
$\mathtt{DLL}(a,\mathbf{nil},c,\mathbf{nil}) \models \exists x,n,b.\ x \mapsto (n,b) * \mathtt{DLL}_{rev}(a,\mathbf{nil},b,x) * \mathtt{DLL}(n,x,c,\mathbf{nil})$	False	2/4	3/5	9/13
$\exists y,a.\ x \mapsto (y,\mathbf{nil}) * y \mapsto (a,x) * \mathtt{DLL}(a,y,c,\mathbf{nil}) \models \mathtt{DLL}(x,\mathbf{nil},c,\mathbf{nil})$	True	3/4	2/4	5/8
$\mathtt{DLL}(x,\mathbf{nil},c,\mathbf{nil}) \models \exists y,a.\ x \mapsto (\mathbf{nil},y) * y \mapsto (a,x) * \mathtt{DLL}(a,y,c,\mathbf{nil})$	False	2/4	3/4	8/10
$\mathtt{TREE}_{pp}(a,\mathbf{nil}) \models \mathtt{TREE}^{rev}_{pp}(a,\mathbf{nil})$	True	2/4	3/8	6/11
$\mathtt{TREE}^{rev}_{pp}(a,\mathbf{nil}) \models \mathtt{TREE}_{pp}(a,\mathbf{nil})$	True	3/8	2/4	5/10
$\mathtt{TLL}_{pp}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{TLL}^{rev}_{pp}(a,\mathbf{nil},c,\mathbf{nil})$	True	4/8	5/10	16/26
$\mathtt{TLL}^{rev}_{pp}(a,\mathbf{nil},c,\mathbf{nil}) \models \mathtt{TLL}_{pp}(a,\mathbf{nil},c,\mathbf{nil})$	True	5/10	4/8	13/22
$\exists l,r,z.\ a \mapsto (l,r,\mathbf{nil},\mathbf{nil}) * \mathtt{TLL}(l,c,z) * \mathtt{TLL}(r,z,\mathbf{nil}) \models \mathtt{TLL}(a,c,\mathbf{nil})$	True	4/7	4/8	13/22
$\mathtt{TLL}(a,c,\mathbf{nil}) \models \exists l,r,z.\ a \mapsto (l,r,\mathbf{nil},\mathbf{nil}) * \mathtt{TLL}(l,c,z) * \mathtt{TLL}(r,z,\mathbf{nil})$	False	4/8	4/7	13/21

$$
\begin{aligned}
\mathtt{DLL}_{rev}(hd,p,tl,n) &\equiv hd \mapsto (n,p) \ \wedge\ hd = tl \mid \exists x.\ tl \mapsto (n,x) * \mathtt{DLL}_{rev}(hd,p,x,tl)\\
\mathtt{DLL}_{mid}(hd,p,tl,n) &\equiv hd \mapsto (n,p) \ \wedge\ hd = tl \mid hd \mapsto (tl,p) * tl \mapsto (n,hd)\\
&\quad \mid \exists x,y,z.\ x \mapsto (y,z) * \mathtt{DLL}(y,x,tl,n) * \mathtt{DLL}_{rev}(hd,p,z,x)\\
\mathtt{TREE}_{pp}(x,b) &\equiv x \mapsto (nil,nil,b) \mid \exists l,r.\ x \mapsto (l,r,b) * \mathtt{TREE}_{pp}(l,x) * \mathtt{TREE}_{pp}(r,x)\\
\mathtt{TREE}^{rev}_{pp}(t,b) &\equiv t \mapsto (nil,nil,b) \mid \exists x,up.x \mapsto (nil,nil,up) * \mathtt{TREE}^{aux}_{pp}(t,b,up,x)\\
\mathtt{TREE}^{aux}_{pp}(t,b,x,d) &\equiv \exists r.x \mapsto (d,r,b) * \mathtt{TREE}_{pp}(r,x) \ \wedge\ x = t\\
&\quad \mid \exists l.x \mapsto (l,d,b) * \mathtt{TREE}_{pp}(l,x) \ \wedge\ x = t\\
&\quad \mid \exists r,up.x \mapsto (d,r,up) * \mathtt{TREE}^{aux}_{pp}(t,b,up,x) * \mathtt{TREE}_{pp}(r,x)\\
&\quad \mid \exists l,up.x \mapsto (l,d,up) * \mathtt{TREE}^{aux}_{pp}(t,b,up,x) * \mathtt{TREE}_{pp}(l,x)\\
\mathtt{TLL}_{pp}(r,p,ll,lr) &\equiv r \mapsto (nil,nil,p,lr) \ \wedge\ r = ll\\
&\quad \mid \exists x,y,z.\ r \mapsto (x,y,p,nil) * \mathtt{TLL}_{pp}(x,r,ll,z) * \mathtt{TLL}_{pp}(y,r,z,lr)\\
\mathtt{TLL}^{rev}_{pp}(t,p,ll,lr) &\equiv ll \mapsto (nil,nil,p,lr) \ \wedge\ ll = t\\
&\quad \mid \exists up,z.ll \mapsto (nil,nil,up,z) * \mathtt{TLL}^{aux}_{pp}(t,p,up,ll,z,lr)\\
\mathtt{TLL}^{aux}_{pp}(t,p,x,d,z,lr) &\equiv \exists r.x \mapsto (d,r,p,nil) * \mathtt{TLL}_{pp}(r,x,z,lr) \ \wedge\ x = t\\
&\quad \mid \exists r,up,q.x \mapsto (d,r,up,nil) * \mathtt{TLL}^{aux}_{pp}(t,p,up,x,q,lr) * \mathtt{TLL}_{pp}(r,x,z,q)
\end{aligned}
$$

Fig. 6. Inductive definitions of predicates used in experiments.

For each experiment, we first replaced each left/right-hand side of an entailment query containing a points-to predicate by a new top-level predicate and then performed the parameter elimination. We have not yet implemented this step, but we tightly followed Alg. 2. Subsequently, we have applied our tool. Column *Answer* of Table 1 shows the results given by the tool. The answers provided by the tool are all correct despite the incompleteness of the approach for non-local systems. The experiments show that our approach is indeed capable of handling quite complex predicates, including the possibility of encoding the same data structures starting from different entry points. Moreover, the approach is also rather efficient: columns $|A_{lhs}|$, $|A_{rhs}|$, and $|A^r_{rhs}|$ provide information about the number of states and transitions of the respective automata. Note that no reduction technique was applied to reduce the size of A^r_{rhs}.

7 Conclusion

We have presented a novel automata-theoretic decision procedure for the entailment problem of a non-trivial subset of SL with inductive predicates. Our reduction to TA can

deal with the problem that the same recursive structure may be represented differently when viewed from different entry points. To deal with this problem, our procedure uses a special closure operation which closes a given TA representation with respect to all rotations of its spanning trees. Our procedure is sound and complete for recursive structures in which all edges are local with respect to a spanning tree. For this case, using our reduction, we show that the entailment problem is EXPTIME-complete. For structures outside of this class, our procedures is incomplete but still sound. We have implemented our approach in a prototype tool which we tested through a number of non-trivial experiments. In the experiments, the incompleteness of our approach has never showed up. Moreover, the experiments also showed that our approach is quite efficient. In the future, we plan to improve the implementation, extend the experiments, as well as to integrate our decision procedure into some verification tool dealing with programs with dynamic linked structures.

References

1. Antonopoulos, T., Gorogiannis, N., Haase, C., Kanovich, M., Ouaknine, J.: Foundations for decision problems in separation logic with general inductive predicates. In: To appear in Proc. of FOSSACS'14 (2014)
2. Bar-Hillel, Y., Perles, M., Shamir, E.: On formal properties of simple phrase structure grammars. Zeitschrift fr Phonetik, Sprachwissenschaft und Kommunikationsforschung 14(2), 143–172 (1961)
3. Berdine, J., Calcagno, C., Cook, B., Distefano, D., O'Hearn, P., Wies, T., Yang, H.: Shape analysis for composite data structures. In: Proc. CAV'07. LNCS, vol. 4590. Springer (2007)
4. Berdine, J., Calcagno, C., O'Hearn, P.W.: A decidable fragment of separation logic. In: Proc. of FSTTCS'04. LNCS, vol. 3328. Springer (2004)
5. Bouajjani, A., Habermehl, P., Holik, L., Touili, T., Vojnar, T.: Antichain-based universality and inclusion testing over nondeterministic finite tree automata. In: Proc. of CIAA. LNCS, vol. 5148. Springer (2008)
6. Brotherston, J., Kanovich, M.: Undecidability of propositional separation logic and its neighbours. In: Proceedings of the 2010 25th Annual IEEE Symposium on Logic in Computer Science. pp. 130–139. LICS '10 (2010)
7. Calcagno, C., Distefano, D.: Infer: An automatic program verifier for memory safety of c programs. In: Proc. of NASA Formal Methods'11. LNCS, vol. 6617. Springer (2011)
8. Comon, H., Dauchet, M., Gilleron, R., Jacquemard, F., Lugiez, D., Tison, S., Tommasi, M.: Tree Automata Techniques and Applications (2005),
   URL: `http://www.grappa.univ-lille3.fr/tata`
9. Cook, B., Haase, C., Ouaknine, J., Parkinson, M.J., Worrell, J.: Tractable reasoning in a fragment of separation logic. In: Proc. of CONCUR'11. LNCS, vol. 6901. Springer (2011)
10. Dudka, K., Peringer, P., Vojnar, T.: Predator: A practical tool for checking manipulation of dynamic data structures using separation logic. In: Proc. of CAV'11. LNCS, vol. 6806. Springer (2011)
11. Enea, C., Saveluc, V., Sighireanu, M.: Compositional invariant checking for overlaid and nested linked lists. In: Proc. of ESOP'13. pp. 129–148 (2013)
12. Iosif, R., Rogalewicz, A., Vojnar, T.: Slide: Separation logic with inductive definitions,
   URL: `http://www.fit.vutbr.cz/research/groups/verifit/tools/slide/`
13. Iosif, R., Rogalewicz, A., Simacek, J.: The tree width of separation logic with recursive definitions. In: Proc. of CADE-24. LNCS, vol. 7898. Springer (2013)

14. Lengal, O., Simacek, J., Vojnar, T.: Vata: a tree automata library,
URL: http://www.fit.vutbr.cz/research/groups/verifit/tools/libvata/
15. Nguyen, H.H., Chin, W.N.: Enhancing program verification with lemmas. In: Proc of CAV'08. LNCS, vol. 5123. Springer (2008)
16. Piskac, R., Wies, T., Zufferey, D.: Automating separation logic using smt. In: Proc. of CAV'13. LNCS, vol. 8044 (2013)
17. Reynolds, J.: Separation Logic: A Logic for Shared Mutable Data Structures. In: Proc. of LICS'02. IEEE CS Press (2002)

Algorithm 2 Elimination of propagated formal parameters

input A rooted system $\langle \mathcal{P}, P_i \rangle$, where $\mathcal{P} = \{P_i(x_{i,1}, \ldots, x_{i,n_i}) \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$, and $P_i(\mathbf{x})$
output A rooted system $\langle Q, Q_i \rangle$, where Q_i has empty formal parameter tuple
global $\mathcal{Q} \leftarrow \mathcal{P}$, `Visited` $\leftarrow \emptyset$, `Parameters` $\leftarrow \emptyset$, `Tracked`

function ELIMINATEPARAMETERS($P_i, \overline{\alpha}$) ▷ $\overline{\alpha}$ is the tuple of actual parameters
  for all $k = 1, \ldots, n_i$ **do** ▷ iterate through all parameters of P_i
    `Tracked` $\leftarrow \overline{\alpha}_k$
    TRACKELIMINATE($P_i, x_{i,k}$, **true**)
    `Parameters` $\leftarrow$ `Parameters` $\cup \{$`Tracked`$\}$
  return TRIM(Q, P_i)

function TRACKELIMINATE($P_r, x_{r,s}$, `del`) ▷ parameter $x_{r,s}$ of P_r, boolean `del`
  $\mathbf{x}_{new} \leftarrow$ **if** `del` **then** $\mathbf{x}_{\neg x_{r,s}}$ **else** $\mathbf{x}$ ▷ initialize a new formal parameter tuple
  $Q_r(\mathbf{x}_{new}) \leftarrow$ **empty_predicate**
  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{Q_r\}$ ▷ create a new predicate name
  for all $q = 1, \ldots, m_r$ **do** ▷ iterate through the rules of P_r
    assume $R_{r,q}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto (\mathbf{y}) * P_{i_1}(\mathbf{u}_1) * \ldots * P_{i_m}(\mathbf{u}_m)$
    if $x_{r,s} \equiv \alpha$ **or** $x_{r,s} \notin \bigcup_{j=1}^{m} \mathbf{u}_j$ **then** ▷ $x_{r,s}$ is allocated in $R_{r,q}$ or not propagated
      $R_{new}(\mathbf{x}_{new}) \leftarrow R_{r,q}[$`Tracked`$/x_{r,s}]$ ▷ replace $x_{r,s}$ in $R_{r,q}$ by the global `Tracked`
    else ▷ check if the current tracked parameter $x_{r,s}$ is passed to P_j on position ℓ
      if $\exists j, \ell \,.\, 1 \leq j \leq m \wedge 0 \leq \ell < |\mathbf{y}_j| \wedge x_{r,s} \equiv (\mathbf{u}_j)_\ell$ **then**
        if the choice of j and ℓ is not unique **then**
          error ("branching propagation for parameter $x_{r,s}$")
        if $x_{r,s} \in \mathbf{y}$ **then** ▷ the tracked parameter is referenced before being passed
          $R_{new}(\mathbf{x}_{new}) \leftarrow$ **if** `del` **then** $\exists x_{r,s} \,.\, R_{r,q}$ **else** $R_{r,q}$
          `del` $\leftarrow$ **false**
        else ▷ the tracked parameter is passed without being referenced
          if `del` **then**
            $R_{new}(\mathbf{x}_{new}) \leftarrow \exists \mathbf{z} \,.\, \alpha \mapsto (\mathbf{y}) * P_{i_1}(\mathbf{u}_1) * \ldots * Q_{i_j}(\mathbf{u}_{\neg x_{r,s}}) * \ldots * P_{i_m}(\mathbf{u}_m)$
          else
            $R_{new}(\mathbf{x}_{new}) \leftarrow R_{r,q}$
        if $x_{j,\ell} \notin$ `Visited` **then** ▷ continue tracking parameter $x_{j,\ell}$ of P_j
          `Visited` $\leftarrow$ `Visited` $\cup \{x_{j,\ell}\}$
          TRACKELIMINATE($P_j, x_{j,\ell}$, `del`)
    $Q_r \equiv Q_r \mid R_{new}$

Algorithm 3 Signature computation

input A rooted system $\langle \mathcal{P}, P_i \rangle$, where $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$
global $\mathtt{sig}_1^{fw}, \mathtt{sig}_1^{bw}, \mathtt{sig}_1^{eq}, \ldots, \mathtt{sig}_n^{fw}, \mathtt{sig}_n^{bw}, \mathtt{sig}_n^{eq}$

function FWEDGE($\mathtt{index}, \mathtt{pred}_{\mathtt{no}}, \mathcal{P}$)
  for all $i = 1, \ldots, n$ **do** ▷ iterate through all predicates
    for all $j = 1, \ldots, m_i$ **do** ▷ iterate through all rules
      assume $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto \mathbf{y} * P_{i_0}(\mathbf{x}_0) * \ldots * P_{i_k}(\mathbf{x}_k) \wedge \Pi$
      $\mathtt{res} \leftarrow$ **true**
      for all $\ell = 0, \ldots, k$ **do** ▷ iterate through all predicate occurrences
        if $i_\ell = \mathtt{pred}_{\mathtt{no}}$ **then**
          if $(\mathbf{x}_\ell)_{\mathtt{index}} \notin \mathbf{y}$ **then** ▷ parameter not referenced in occurrence of $P_{\mathtt{pred}_{no}}$
            $\mathtt{res} \leftarrow$ **false**
  return res

function BWEDGE($\mathtt{index}, \mathtt{pred}_{\mathtt{no}}, \mathcal{P}$)
  for all $i = 1, \ldots, n$ **do** ▷ iterate through all predicates
    for all $j = 1, \ldots, m_i$ **do** ▷ iterate through all rules
      assume $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto \mathbf{y} * P_{i_0}(\mathbf{x}_0) * \ldots * P_{i_k}(\mathbf{x}_k) \wedge \Pi$
      $\mathtt{res} \leftarrow$ **true**
      for all $\ell = 0, \ldots, k$ **do** ▷ iterate through all predicate occurrences
        if $i_\ell = \mathtt{pred}_{\mathtt{no}}$ **then**
          if $(\mathbf{x}_\ell)_{\mathtt{index}} \neq_\Pi \alpha$ **then** ▷ parameter not allocated in occurrence of $P_{\mathtt{pred}_{no}}$
            $\mathtt{res} \leftarrow$ **false**
  return res

function ALLOCATED(index,$\mathtt{pred}_{\mathtt{no}}$,$\mathcal{P}$)
  $\mathtt{res} \leftarrow$ **true**
  for all $j = 1, \ldots, m_{\mathtt{pred}_{\mathtt{no}}}$ **do** ▷ iterate through all rules of $P_{\mathtt{pred}_{\mathtt{no}}}$
    assume $R_{\mathtt{pred}_{\mathtt{no}},j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto \mathbf{y} * P_{i_0}(\mathbf{x}_0) * \ldots * P_{i_k}(\mathbf{x}_k) \wedge \Pi$
    if $(\mathbf{x})_{\mathtt{index}} \neq_\Pi \alpha$ **then** ▷ formal parameter index not allocated in $P_{\mathtt{pred}_{\mathtt{no}}}$
      $\mathtt{res} \leftarrow$ **false**
  return res

function REFERENCED(index,$\mathtt{pred}_{\mathtt{no}}$,$\mathcal{P}$)
  $\mathtt{res} \leftarrow$ **true**
  for all $j = 1, \ldots, m_{\mathtt{pred}_{\mathtt{no}}}$ **do** ▷ iterate through all rules of $P_{\mathtt{pred}_{\mathtt{no}}}$
    assume $R_{\mathtt{pred}_{\mathtt{no}},j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto \mathbf{y} * P_{i_0}(\mathbf{x}_0) * \ldots * P_{i_k}(\mathbf{x}_k) \wedge \Pi$
    if $(\mathbf{x})_{\mathtt{index}} \notin \mathbf{y}$ **then** ▷ formal parameter index not referenced in $P_{\mathtt{pred}_{\mathtt{no}}}$
      $\mathtt{res} \leftarrow$ **false**
  return res

function COMPUTESIGNATURES($\mathcal{P}$)
  for all $i = 1, \ldots, n$ **do** ▷ iterate through all predicates of $\mathcal{P}$
    assume $P_i(\mathbf{x})$ ▷ P_i has formal parameters $\mathbf{x}$
    for all $p = 0, \ldots, |\mathbf{x}| - 1$ **do** ▷ iterate through all formal parameters of P_i
      $\mathtt{sig}_i^{fw} \leftarrow \emptyset, \mathtt{sig}_i^{bw} \leftarrow \emptyset, \mathtt{sig}_i^{eq} \leftarrow \emptyset$
      if FWEDGE$(i, p, \mathcal{P})$ **and** ALLOCATED$(i, p, \mathcal{P})$ **then**
        $\mathtt{sig}_i^{fw} \leftarrow \mathtt{sig}_i^{fw} \cup \{p\}$
      else
        if BWEDGE$(i, p, \mathcal{P})$ **and** REFERENCED$(i, p, \mathcal{P})$ **then**
          $\mathtt{sig}_i^{bw} \leftarrow \mathtt{sig}_i^{bw} \cup \{p\}$
        else
          $\mathtt{sig}_i^{eq} \leftarrow \mathtt{sig}_i^{eq} \cup \{p\}$

Algorithm 4 Converting rooted inductive systems to tree automata

input A rooted system $\langle \mathcal{P}, P_{\mathtt{index}} \rangle$, $\mathcal{P} = \{P_i \equiv |_{j=1}^{m_i} R_{i,j}\}_{i=1}^{n}$, and actual parameters $\mathbf{u}$ of a call of $P_{\mathtt{index}}$
output A tree automaton $A = \langle \Sigma, Q, \Delta, F \rangle$

function SL2TA($\mathcal{P}, \mathtt{index}, \mathbf{u}$)
  $\mathcal{P} \leftarrow$ SPLITSYSTEM($\mathcal{P}$)
  $\mathcal{P} \leftarrow$ ELIMINATEPARAMETERS($P_{\mathtt{index}}, \mathbf{u}$)
  COMPUTESIGNATURES($\mathcal{P}$)
  $\Delta \leftarrow \emptyset, \Sigma \leftarrow \emptyset$
  for all $i = 1, \ldots, n$ **do** ▷ iterate through all predicates of $\mathcal{P}$
    for all $j = 1, \ldots, m_i$ **do** ▷ iterate through all rules of P_i
      assume $R_{i,j}(\mathbf{x}) \equiv \exists \mathbf{z} \,.\, \alpha \mapsto \overline{\beta} * P_{i_0}(\mathbf{z}_0) * \ldots * P_{i_k}(\mathbf{z}_k) \wedge \Pi$
      $\phi \leftarrow$ **if** $\alpha \in \mathtt{Parameters}$ **then** $\alpha \mapsto (\overline{\beta}) \wedge \Pi$ **else** $\exists \alpha' \,.\, \alpha' \mapsto (\overline{\beta}) \wedge \Pi \wedge \alpha = \alpha'$
      $\mathbf{x}_{-1}^{fw} \leftarrow \langle\rangle, \mathbf{x}_{-1}^{bw} \leftarrow \langle\rangle, \mathbf{x}_{-1}^{eq} \leftarrow \langle\rangle$ ▷ initialize input ports
      for all $\ell = 0, \ldots, |\mathbf{x}| - 1$ **do** ▷ iterate through the formal parameters of $R_{i,j}$
        if $\ell \in \mathtt{sig}_i^{fw}$ **then**
          $\mathbf{x}_{-1}^{fw} \leftarrow \mathbf{x}_{-1}^{fw} \cdot \langle (\mathbf{x})_\ell \rangle$
        if $\ell \in \mathtt{sig}_i^{bw}$ **then**
          $\mathbf{x}_{-1}^{bw} \leftarrow \mathbf{x}_{-1}^{bw} \cdot \langle (\mathbf{x})_\ell \rangle$
        if $\ell \in \mathtt{sig}_i^{eq}$ **then**
          $\mathbf{x}_{-1}^{eq} \leftarrow \mathbf{x}_{-1}^{eq} \cdot \langle (\mathbf{x})_\ell \rangle$
      SELECTORSORT($\mathbf{x}_{-1}^{bw}, \overline{\beta}$)
      for all $\ell = 0, \ldots, k$ **do** ▷ iterate through predicate occurrences in $R_{i,j}$
        $\mathbf{x}_\ell^{fw} \leftarrow \langle\rangle, \mathbf{x}_\ell^{bw} \leftarrow \langle\rangle, \mathbf{x}_\ell^{eq} \leftarrow \langle\rangle$ ▷ initialize output ports
        for all $r = 0, \ldots, |\mathbf{z}_\ell| - 1$ **do** ▷ iterate through variables of occurrence $P_{i_\ell}(\mathbf{z}_\ell)$
          if $r \in \mathtt{sig}_{i_\ell}^{fw}$ **then**
            $\mathbf{x}_\ell^{fw} \leftarrow \mathbf{x}_\ell^{fw} \cdot \langle (\mathbf{z}_\ell)_r$
          if $r \in \mathtt{sig}_{i_\ell}^{bw}$ **then**
            $\mathbf{x}_\ell^{bw} \leftarrow \mathbf{x}_\ell^{bw} \cdot \langle (\mathbf{z}_\ell)_r \rangle$
          if $r \in \mathtt{sig}_{i_\ell}^{eq}$ **then**
            $\mathbf{x}_\ell^{eq} \leftarrow \mathbf{x}_\ell^{eq} \cdot \langle (\mathbf{z}_\ell)_r \rangle$
        SELECTORSORT($\mathbf{x}_\ell^{fw}, \overline{\beta}$)
      $\mathtt{T} \leftarrow \langle \phi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_k \rangle$ ▷ create a new tile
      $\mathtt{lhs} \leftarrow \langle q_{i_0}, \ldots, q_{i_k} \rangle$ ▷ create the left hand side of the transition rule for $R_{i,j}$
      $(\mathtt{T}_{\mathtt{new}}, \mathtt{lhs}_{\mathtt{new}}) \leftarrow$ SORTTILE($\mathtt{T}, \mathtt{lhs}$)
      $\Sigma \leftarrow \Sigma \cup \{\overline{\mathtt{T}_{\mathtt{new}}}\}$ ▷ insert normalized tile in the alphabet
      $Q \leftarrow Q \cup \{q_i, q_{i_0}, \ldots, q_{i_k}\}$
      $\Delta \leftarrow \Delta \cup \{\overline{\mathtt{T}_{\mathtt{new}}}(\mathtt{lhs}_{\mathtt{new}}) \to q_i\}$ ▷ build transition using normalized tile
  $F = \{q_{\mathtt{index}}\}$
  return $A = \langle Q, \Sigma, \Delta, F \rangle$

Algorithm 5 Rotation Closure of Quasi-canonical TA

input A quasi-canonical tree automaton $A = \langle Q, \Sigma, \Delta, F \rangle$
output A tree automaton A^r, where $\mathcal{L}(A^r) = \{u : \mathbb{N}^* \rightharpoonup_{fin} \mathcal{T}^{qc} \mid \exists t \in \mathcal{L}(A) \,.\, u \sim^{qc} t\}$

$A^r = \langle Q_r, \Sigma, \Delta_r, F_r \rangle \leftarrow A$ ▷ make a copy of A into A^r
for all $\rho \in \Delta$ **do** ▷ iterate through the rules of A
  assume $\rho \equiv T(q_0, \ldots, q_k) \to q$ **and** $T \equiv \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_k \rangle$ ▷ the chosen rule ρ will recognize the root of the rotated tree
  if $\mathbf{x}_{-1} \neq \emptyset$ **or** $q \notin F$ **then**
    assume $\mathbf{x}_{-1} \equiv \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{eq}$ ▷ factorize the input port
    if $\mathbf{x}_{-1}^{bw} \neq \emptyset$ **then**
      $Q^{rev} \leftarrow \{q^{rev} \mid q \in Q\}$ ▷ states q^{rev} label the unique reversed path in A^r
      $(Q_\rho, \Delta_\rho) \leftarrow (Q \cup Q^{rev} \cup \{q_\rho^f\}, \Delta)$ ▷ assuming $Q \cap Q^{rev} = \emptyset$, $q_\rho^f \notin Q \cup Q^{rev}$
      $p \leftarrow \text{POSITIONOF}(\mathbf{x}_{-1}^{bw}, \varphi)$ ▷ find new output port for $\mathbf{x}_{-1}$ based on selectors of φ
      $T_{new} \leftarrow \langle \varphi, \emptyset, \mathbf{x}_0, \ldots, \mathbf{x}_p, \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{eq}, \mathbf{x}_{p+1}, \ldots, \mathbf{x}_k \rangle$ ▷ swap $\mathbf{x}_{-1}^{bw}$ with $\mathbf{x}_{-1}^{fw}$
      $\Delta_\rho \leftarrow \Delta_\rho \cup \{T_{new}(q_0, \ldots, q_p, q^{rev}, q_{p+1}, \ldots, q_k) \to q_\rho^f\}$
      $(\Delta_\rho, _) \leftarrow \text{ROTATERULE}(q, \Delta, \Delta_\rho, \emptyset, F)$ ▷ continue building Δ_ρ recursively
      $A_\rho \leftarrow \langle Q_\rho, \Sigma, \Delta_\rho, \{q_\rho^f\} \rangle$ ▷ A_ρ recognizes trees whose roots are labeled by ρ
      $A^r \leftarrow A^r \cup A_\delta$ ▷ incorporate A_δ into A_r
return A^r

function ROTATERULE(q, Δ, Δ_{new}, `visited`, F)
  `visited` $\leftarrow$ `visited` $\cup \{q\}$
  for all $(U(s_0, \ldots, s_\ell) \to s) \in \Delta$ **do** ▷ iterate through the rules of A
    for all $0 \leq j \leq \ell$ **such that** $s_j = q$ **do** ▷ for all occurrences of q on the rhs
      assume $U = \langle \varphi, \mathbf{x}_{-1}, \mathbf{x}_0, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_\ell \rangle$
      assume $\mathbf{x}_j \equiv \mathbf{x}_j^{fw} \cdot \mathbf{x}_j^{bw} \cdot \mathbf{x}_j^{eq}$ ▷ factorize the $\mathbf{x}_j$ output port
      if $\mathbf{x}_{-1} = \emptyset$ **and** $s \in F$ **then** ▷ remove $\mathbf{x}_j$ from output and place it as input port
        $U' \leftarrow \langle \varphi, \mathbf{x}_j^{bw} \cdot \mathbf{x}_j^{fw} \cdot \mathbf{x}_j^{eq}, \mathbf{x}_0, \ldots, \mathbf{x}_{j-1}, \mathbf{x}_{j+1}, \ldots, \mathbf{x}_\ell \rangle$
        $\Delta_{new} \leftarrow \Delta_{new} \cup \{U'(s_0, \ldots, s_{j-1}, s_{j+1}, \ldots, s_\ell) \to q^{rev}\}$
      else ▷ swap $\mathbf{x}_{-1}$ with $\mathbf{x}_j$
        $\mathbf{x}_{-1} \equiv \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{eq}$
        if $\mathbf{x}_{-1}^{bw} \neq \emptyset$ **then**
          `ports` $\leftarrow \langle \mathbf{x}_0, \ldots, \mathbf{x}_{j-1}, \mathbf{x}_{j+1}, \ldots, \mathbf{x}_\ell \rangle$
          `states` $\leftarrow (s_0, \ldots, s_{j-1}, s_{j+1}, \ldots, s_\ell)$
          $p \leftarrow \text{INSERTOUTPORT}(\mathbf{x}_{-1}^{bw} \cdot \mathbf{x}_{-1}^{fw} \cdot \mathbf{x}_{-1}^{eq}$, `ports`, $\varphi)$
          INSERTLHSSTATE(s^{rev}, `states`, p)
          $U_{new} \leftarrow \langle \varphi, \mathbf{x}_j^{bw} \cdot \mathbf{x}_j^{fw} \cdot \mathbf{x}_j^{eq}$, `ports` $\rangle$ ▷ create rotated tile
          $\Delta_{new} \leftarrow \Delta_{new} \cup \{U_{new}($`states`$) \to q^{rev}\}$ ▷ create rotated rule
          if $s \notin$ `visited` **then**
            $(\Delta_{new},$ `visited`$) \leftarrow \text{ROTATERULE}(s, \Delta, \Delta_{new},$ `visited`$, F)$
  return $(\Delta_{new},$ `visited`$)$